



D4.4

Overall system validation and assessment

Project number:	619086
Project acronym:	MAMMOET
Project title:	Massive MIMO for Efficient Transmission
Project Start Date:	1 January, 2014
Duration:	36 months
Programme:	FP7/2007-2013
Deliverable Type:	Report
Reference Number:	ICT-619086-D4.4
Workpackage:	WP 4
Due Date:	31 December, 2016
Actual Submission Date:	16 January, 2017
Responsible Organisation:	IMEC
Editor:	Claude Desset
Dissemination Level:	PU
Revision:	1.0
Abstract:	Over the course of the Mammoet project, different aspects of the MaMi technology have been investigated, including theoretical studies, simulations, hardware design and testbed-based experiments. This deliverable validates aspects connecting the work of multiple workpackages and project components. This includes system simulations based on the channel model derived from measurements, as well as cross-validation of power consumption figures with respect to power models.
Keywords:	Massive MIMO, channel model, system simulations, power consumption



This project has received funding from the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement no. 619086.

Editor

Claude Desset (IMEC)

Contributors (ordered according to beneficiary numbers)

Claude Desset (IMEC)

Liesbet Van der Perre (KUL)

Liang Liu, Ove Edvors (ULUND)

Javier Lorca Hernando (TID)

Executive Summary

This deliverable performs simulation-based overall system validation and integrates actual MAMMOET results, solutions and measurements from the different workpackages. It focuses on the most relevant target scenarios where MaMi can bring a strong improvement in spectral and power efficiency. Those scenarios are tuned based on the latest findings: a 100×10 MaMi configuration is proposed as replacement of macro base stations, using for best performance ZF precoding and frequency-domain smoothing of the channel estimates.

The specific channel model developed in MAMMOET is integrated into the system simulator. The properties of the corresponding channels are analyzed with a special focus on its impact on the overall performance. A few tests were also performed based on actual channel traces extracted from the testbed. While good performance is obtained in all cases, some situations lead to an increased degradation as compared to theoretical model. This is especially the case when some users are closely located, when increasing the number of users or when uncompensated path loss differences are present between users. In such cases specific measures may be needed, such as limiting the modulation order, reducing the number of users, or investigating more advanced scheduling and power control options. Precoder normalization based on hardware or regulation constraints leads to a limited degradation of the performance.

The overall performance is simulated over the selected MAMMOET scenario and solution. The simulation includes models of non-idealities representing the impact of the new modulator design (digital transmitter from workpackage 2) and validating its quality: the measured distortion at -45 dBc is much below the value which would lead a visible degradation in MaMi operation.

The overall power consumption is re-assessed and compared to traditional designs. It is also cross-validated against hardware design measurements. A good match is observed between the model assumptions and the actual hardware components. More importantly, the overall power assessment over the selected scenario confirms a strong benefit from MaMi as compared to traditional architectures: at least a 10-fold power reduction for the same spectral efficiency, and even more with advanced architectures.

Contents

1	Introduction	1
2	Simulation scenarios for validation	2
2.1	Application scenarios	2
2.2	Simulation scenarios	2
3	Validation simulations with measurement-based Massive MIMO channel model	5
3.1	Introduction	5
3.2	Checking normalization options for different channels/precoders	6
3.2.1	Rayleigh	6
3.2.2	Lund measurements (un-normalized users)	7
3.2.3	Lund measurements (normalized users)	7
3.3	Detailed channel analysis	11
3.4	Conclusions	14
4	Simulation-based performance validation incorporating MAMMOET solutions	16
4.1	Overall performance	16
4.2	Comparison to the testbed performance	17
4.2.1	Testbed scenarios	17
4.2.2	Cross-validation between testbed and simulations	17
4.3	Simulated performance with hardware non-idealities	18
5	Overall power consumption including hardware components	20
5.1	Power consumption of digital baseband hardware	20
5.2	Overall power consumption	20
6	Conclusions	25
7	Abbreviations	26

List of Figures

3.1	Impact of different normalizations on MRT performance for Rayleigh multi-tap channels.	7
3.2	Impact of different normalizations on ZF performance for Rayleigh multi-tap channels.	8
3.3	Impact of different normalizations on MRT performance for Lund channels. . . .	8
3.4	Impact of different normalizations on ZF performance for Lund channels. . . .	9
3.5	Impact of different normalizations on MRT performance for Lund channels normalized over users.	10
3.6	Impact of different normalizations on ZF performance for Lund channels normalized over users.	10
3.7	Normalization coefficient spread after averaging over different dimensions (Rayleigh channel).	11
3.8	Normalization coefficient spread after averaging over different dimensions (normalized Lund channel).	12
3.9	Statistics of a typical Rayleigh channel H (100×9 , 72 i.i.d. taps, frequency analysis over 1200 subcarriers): time-domain delay profile (upper left), frequency-domain response (lower left), cross-correlation between users (upper right), condition number of HH^H (lower right 1–1200) and total energy of the 100 antennas (lower-right 1201–1300).	12
3.10	Statistics of a measurement-based Lund channel (100×9) after user normalization.	13
3.11	Statistics of a measurement-based Lund channel (100×9) without user normalization.	14
4.1	MaMi performance for the scenario of Section 2.2 with real MaMi channel model and estimated CSI, compared to scenarios with AWGN channels and/or ideal CSI.	17
4.2	MaMi performance when using a channel trace from the testbed, as compared to either AWGN or to the MaMi channel model.	18
4.3	MaMi performance considering the impact of the digital-RF modulator (-45 dBc EVM) as well as strongly degraded versions of it, as compared to the ideal case.	19
5.1	Block diagram of OFDM-based massive MIMO digital baseband processing, with implemented function blocks highlighted	21
5.2	Power model assessment between macro (config. 1), MaMi with traditional hardware (config. 2) and MaMi with advanced digital-RF (config. 3).	22
5.3	Power split over the 3 phases when using traditional hardware.	24
5.4	Power split over the 3 phases when using advanced digital RF.	24



List of Tables

2.1	Parameters for the main validation scenario.	4
5.1	Digital baseband processing blocks implementation result with ST 28nm FD-SOI	20
5.2	Power modeling parameters differing between macro and MaMi base stations. All parameters are given per sector unless specified.	22
5.3	Detailed power consumption for one sector of the MaMi baseline scenario, con- sidering both analog architecture types.	23

Chapter 1

Introduction

At the end of the MAMMOET project, a number of final validation steps are taken in order to assess the overall potential of the Massive MIMO technology investigated in the project. Those final steps relate to the achievements of the different workpackages in the MAMMOET project, putting together different elements or cross-validating conclusions obtained through different approaches.

Chapter 2 summarizes application and simulation scenarios. Based on the work in Deliverables D1.1 and D4.1, the main scenarios were already described but specific updates are proposed based on the latest research results such as the introduction of channel interpolation investigated in D3.2 and D3.3. Additional results confirm the worse performance of the MRT precoder, hence the use of ZF is recommended as default configuration.

Chapter 3 summarizes the integration of the channel model developed in Deliverables D1.2 and D1.3 into the Matlab simulator of D4.1. This fills an important hole in the system validation as it was previously only possible on theoretical models such as multi-tap Rayleigh, less suited to the specific Massive MIMO environment. Two aspects are considered. The first one is a comparison of different statistical properties of this channel model to the theoretical cases. A second one is a comparison of the system performance when using both models under different precoder normalization constraints.

Chapter 4 presents simulation-based validation for the selected baseline scenario, a comparison to testbed scenarios and an investigation on the performance impact of the signal accuracy level reported from the digital-RF modulator hardware (D4.3). The impact of channel estimation and the different channels is also further explored.

Chapter 5 updates power consumption estimates and cross-validates them with hardware measurements. As compared to power modeling in Deliverable D3.2, a number of power modeling assumptions are revisited based on latest research results, especially related to the impact of channel estimation and interpolation. Power consumption values are also updated based on latest technology trends. Results from the Workpackage 2 modulator and its measured power consumption are related and compared to the power modeling trends. The power consumption of an implementation of the most critical digital baseband functions is also reported.

Chapter 2

Simulation scenarios for validation

This chapter reviews and updates scenarios proposed in earlier deliverables (D1.1 [2] and D4.1 [6]), in order to identify the most relevant cases to focus on for the final simulation-based validation.

2.1 Application scenarios

From the operator point of view, the open exhibition scenario described in Deliverable D1.1 is the most attractive scenario where Massive MIMO deployment is expected to make the difference. It targets the outdoor deployment of macro base stations serving a high number of users in a place such as a conference center, theater or stadium. The prevalence of such scenarios is also confirmed by the decision of 3GPP to prioritize the work around enhanced Mobile BroadBand (eMBB) use cases for the first release of NR (5G).

Many of the generic parameters were already taken from Deliverable D1.1 and summarized in Deliverable D4.1 for simulation purposes. Those include the 20 MHz bandwidth, use of OFDM and TDD, LTE-based time and frequency parameters, full-buffer simulations with flexible modulation and coding schemes, etc. Some other parameters specified for D1.1 scenarios are not directly applicable for single point-to-multipoint baseband link simulation but rather for network simulations which are out of the scope of this project. This is the case for, e.g., large-scale geometry and propagation aspects in the cell. Other parameters such as individual antenna gain and receiver noise figure relate more to the link budget computation than to the baseband SNR vs. BER simulation.

2.2 Simulation scenarios

Simulation scenarios have been proposed in Chapter 4 of Deliverable D4.1, in line with application scenarios of Deliverable D1.1. At the end of the MAMMOET project, those scenarios still form a relevant baseline for the final validation, but they are updated by implementing some modifications based on findings in the project.

The most important point is the type of precoder. While MRT/MRC was initially pushed in the Massive MIMO community thanks to its simplicity and ease of theoretical analysis, it is not sufficiently reliable in some deployment scenarios. MRT only performs fine when the system load is very low (considering both the ratio between number of users and number of base station antennas as well as the modulation order). At medium load its performance strongly degrades and the simulated BER starts flooring. Chapter 3 of Deliverable D3.2 [5] has modeled the

specific impairments affecting the MRT performance when increasing the load, supporting this conclusion, which is also confirmed by comparisons in Chapter 3 of the current deliverable.

Another element was the fear of a too high complexity for other linear precoders/detectors such as ZF and MMSE. However, the complexity of the precoder computation and more generally the digital signal processing was found to be well under control, as presented in Deliverable D3.2. ZF is hence selected as default in the final validation. MMSE could be slightly better in some very specific scenarios (such as $K \approx M$, large path loss differences included or multi-cell interference scenarios) but generally offers the same performance as ZF. In downlink the ZF precoder offers the advantage of not having to estimate the noise variance at the transmitter side. In uplink the noise variance can be measured on received signals and can be used to compute an MMSE detector with negligible complexity increase. The limited difference between ZF and MMSE performance comes as another benefit from the channel hardening effect. Indeed, as soon as the number of antennas is slightly larger than the number of users, a well-conditioned matrix is generated and it does not require additional regularization in order to be inverted.

Generally speaking, the following typical parameter ranges are most realistic for Massive MIMO and a subset of them can be used in the final validation:

- $M = 50, 100, 200$
- $K = 1, 5, 10, 20, 40$ (depending on M)
- Modulations: QPSK and 16-QAM with different coding rates

Those parameter ranges fit with what was already proposed in D1.1 and D4.1, i.e., $M = 100$ and $K = 10$ at most. Those are kept as default values.

Another important aspect is channel model and channel estimation strategy. The measurement-based channel model developed at Lund is integrated into the simulator and included in simulations, instead of the theoretical multi-tap Rayleigh model initially used. Concerning channel estimation, the frequency-domain channel interpolation strategy presented and tested in Deliverables D3.2 and D3.3 [7] is included in order to improve the channel estimates as compared to the baseline per-subcarrier solution initially proposed.

The smoothing can be parameterized between more noise suppression and more channel selectivity, as it is based on the maximum delay profile allowed for the channel. Simulations are performed by constraining the duration to half the cyclic prefix or 72 samples, which is already providing a large margin as compared to channels produced by the model (see Figures in Chapter 3, such as 3.10 and 3.11 reporting a delay spread of some 15 samples). More noise suppression could hence be obtained by further reducing the constrained duration but this may not be applicable to other (NLOS) scenarios.

Results can be compared to theoretical predictions in terms of SNR required in order to achieved a target BER such as 10^{-5} . Table 2.1 summarizes all parameters relevant to the selected scenario. While simulations are only performed in quasi-static mode (constant channel per frame), mobility tests have been run using a similar frame using the testbed, with successful performance observed at 50 km/h (see Deliverable D4.2 [8]).

Table 2.1: Parameters for the main validation scenario.

Category	Parameter	Value	Notes
Air interface	Mode	OFDM, TDD	
	Carrier	2 – 3 GHz (LTE-based)	
	Bandwidth	20 MHz	
	FFT size	2048	at 30.72 MHz
	Data subcarriers	1200	(15 kHz spacing)
	Cyclic Prefix (CP)	144	
	Modulation	16-QAM	(or QPSK)
	Coding rate	3/4	LDPC
	Frame (total, training)	(14, 1)	(OFDM symbols)
Massive MIMO	Antennas (M)	100	
	Users (K)	10	
	Precoder	ZF	
	Channel estimation	Per-subcarrier	K sequences
	Channel smoothing	FFT-based	half-CP constraint
	Hardware	ideal	
Link	Reference coverage	macro 49 dBm	(4 × 4)
	Additional margin	3 dB	for MaMi channel estim.
	Channel model	Lund model	measurement-based
	Mobility	Quasi-static	(constant per frame)

Chapter 3

Validation simulations with measurement-based Massive MIMO channel model

3.1 Introduction

This section compares channel statistics and system performance using Rayleigh multi-tap versus Lund measurement-based channel models described in Deliverables D1.2 [3] and D1.3 [4]. Given the very different statistics between both models, the study also covers aspects of precoder normalization, especially the impact of constraining the precoder to have the same total energy over different antennas, users or subcarriers. This is important for practical implementation with respect to having identical power amplifiers, ensuring fairness, and generating a frequency-flat spectrum, respectively.

Indeed, OFDM-based MaMi transmissions rely on a 3D precoder matrix over antennas, users and subcarriers. A global normalization constraint is always applied to the overall precoder energy, after accumulating it over the 3 dimensions. More constraining options are possible at the level of antennas (equal output per PA), users (fairness in absence of power control) and subcarriers (frequency-flat overall output), leading in total to 8 combinations of those 3 binary constraints (antenna-level, user-level and subcarrier-level).

The main objective for this normalization study is to investigate the system behavior combining different channel models and different normalization options and recommend a normalization scheme to be used for further studies and actual deployment. Perfect CSI is assumed in these simulations. The general recommendation concerning precoder choice is clearly ZF instead of MRT but both results are kept for better understanding of the MaMi behavior. When using MRT the normalization constraints can be interpreted in the following way:

- The un-normalized MRT implementation simply considers the conjugate channel elements as precoding coefficients, up to the overall normalization factor.
- On the antenna dimension, this is optimum in the sense that more powerful channel coefficients benefit from the corresponding higher precoding coefficients, in order to reach the largest SNR after accumulation over all antennas, by definition of the matched-filtering property.
- On the user dimension, the un-normalized MRT will tend to optimize system capacity while not providing fairness. The combination with some kind of power control should be

investigated. The expected effect is limited in the performed simulations, as no specific user path loss differences are introduced. In actual deployments large path loss differences are present but they require other compensation techniques that precoder normalization, which are out of the scope of this study (power control, link adaptation, scheduling...).

- On the subcarrier dimension, the un-normalized MRT will lead to increased differences due to inverse waterfilling, i.e., strong subcarriers will be enhanced and have perfect symbols while weak subcarriers will become weaker and dominate the error rate, at least for the uncoded case. The situation for coded cases is however less clear, given that good symbols benefiting from a higher SNR enable to correct erroneous ones. This effect should also be limited when sufficiently many independent antennas are present to average the signals, particularly when taking channel hardening into account.

When using ZF precoding, an additional specific problem is that enforcing antenna normalization generally destroys the perfect interference cancellation property of the precoder. User and subcarrier normalization do not cause this problem.

- The un-normalized ZF precoder in the antenna dimension is optimum in the sense that it ensures perfect interference cancellation
- On the user dimension, unlike the MRT case, the un-normalized ZF precoder will tend to optimize for fairness and not system capacity. This should again be further investigated in line with power control.
- On the subcarrier dimension, it will also tend to compensate bad subcarriers and hence - in the uncoded case - tend to minimize the BER. This may again not be the case with channel coding.

3.2 Checking normalization options for different channels/precoders

This section simulates the 8 possible normalization options over Rayleigh multi-tap, raw Lund channel (including path loss differences among users) and normalized Lund channel (providing the same path loss per user). Normalizing the channel path loss over users is independent of normalizing the precoder on the user dimension.

3.2.1 Rayleigh

Figures 3.1 and 3.2 illustrate MRT and ZF performance for 8 possible normalizations on a Rayleigh multi-tap channel. Symbols A, U and S refer to constraining on antennas, users and subcarriers, respectively. Both figures show that the different normalization options have a similar performance except when all 3 dimensions are constrained together (leading all unit-magnitude components in the 3D precoder matrix), where 1.5 dB is lost. The interpretation is that diversity comes from antennas, users and subcarriers, and whenever at least one dimension remains unconstrained, signals are sufficiently stable in amplitude based on adding/averaging on the remaining dimension(s), which implies that except for the fully-constrained case, the applied normalization constraints have a very limited effect. Most likely only antenna normalization due to hardware constraints would be enforced in practical deployments.




Figure 3.1: Impact of different normalizations on MRT performance for Rayleigh multi-tap channels.

The difference in performance between MRT and ZF is limited to a bit more than 1 dB. This complies with expectations given that the selected QPSK 3/4 has low requirements on SINR (5.5 dB) while the 100×9 configuration provides an interference attenuation by 11 dB, which theoretically implies a degradation by 1.4 dB due to MRT inter-user interference.

3.2.2 Lund measurements (un-normalized users)

Figures 3.3 and 3.4 illustrate MRT and ZF performance for 8 possible normalizations on the measurement-based channel model from Lund university. The performance is generally poor but this partly comes from differences in total channel energy over the different users, given that the average BER is dominated by the user suffering from the worst path loss, in absence of power control. In order to focus only on the impact of the different normalization constraints on different channels assuming equal path loss users, the study is continued in the next sections assuming the Lund channel is renormalized in order to have the same total energy over the different users. Hence the results from this section on un-normalized Lund channels are not further investigated. Generally speaking, path loss differences are an important aspect of cellular communications but they have to be tackled by a combination of power control and link adaptation, i.e., reducing the modulation order and coding rate for the weakest users.

3.2.3 Lund measurements (normalized users)

Figures 3.5 and 3.6 illustrate MRT and ZF performance for 8 possible normalizations on the measurement-based channel model from Lund university, after renormalizing all users to the same total energy. As compared to the un-normalized channel case, the ZF-based optimal




Figure 3.2: Impact of different normalizations on ZF performance for Rayleigh multi-tap channels.




Figure 3.3: Impact of different normalizations on MRT performance for Lund channels.




Figure 3.4: Impact of different normalizations on ZF performance for Lund channels.

performance is strongly improved (by 8 dB) due to the removal of energy differences between user channels. The achieved performance (SNR requirement of -14 dB) is comparable to the Rayleigh case, where users also have similar energy.

The poor performance of MRT likely comes from the stronger correlation between specific users. Indeed, for a 100×9 configuration, the typical correlation between any 2 users is around 0.1 for the Rayleigh channel. With the Lund channel this is found to increase to around 0.5 for specific user pairs (see Section 3.3), which implies that those users will have an SINR already limited to 6 dB by considering only interference from the other user of the pair, and further reduced from additional interference and noise terms. On the contrary, the ZF precoder can cancel this interference, at the cost of some gain reduction especially if the matrix is too strongly ill-conditioned. Channel statistics are further analyzed in Section 3.3.

Concerning normalization, a significant degradation is observed for modes combining antenna and user normalization, while subcarrier normalization has limited impact or is actually slightly positive. A likely explanation is that some (antenna, user) combinations have much lower energy, even after averaging over the subcarrier dimension. This probably comes from geometrical parameters, such as the use of the cylindric antenna array and possibly no signal for users located exactly at the other side of the cylinder for the selected antenna. Hence, when trying to normalize on both antennas and users, the corresponding positions require a large change in coefficient energy which degrades the orthogonality between users. Figures 3.7 and 3.8 illustrate this phenomenon by comparing the distribution of normalization coefficients for Rayleigh and normalized Lund channels and for the 8 possible normalization constraints. When all coefficients are close to one, the normalization constraint will have a limited or no effect on the precoder, while when the range is wide, the normalization stage will strongly impact the precoder. On Rayleigh only the most constrained "A U S" case has a wide distribution and corresponds to a significant degradation as those widely-distributed coefficients multiply




Figure 3.5: Impact of different normalizations on MRT performance for Lund channels normalized over users.




Figure 3.6: Impact of different normalizations on ZF performance for Lund channels normalized over users.




Figure 3.7: Normalization coefficient spread after averaging over different dimensions (Rayleigh channel).

different precoder coefficients. On normalized Lund channels both "A U -" and "A U S" show a wide distribution, which matches the observed performance degradation for those 2 modes. Future research should validate those results for channel models based on planar arrays instead of cylindrical array, as actual deployments are expected to build on rectangular arrays.

3.3 Detailed channel analysis

In order to further understand the impact of the different channels, a number of plots of channel data and derived statistics are presented. Figure 3.9 illustrates a Rayleigh channel. It is created as a 3D matrix of 100 (antennas) $\times$ 9 (users) $\times$ 72 (taps) of equal energy i.i.d. complex Gaussian random variables. The upper left sub-figure illustrates the time-domain channel traces for all 900 combinations (100 antennas and 9 users), showing the 72 taps of equal expected energy. The black curve in the middle around amplitude 0.12 corresponds to the averaged amplitude over all those (antenna, user) combinations, with an expected value $1/\sqrt{72}$. The lower left figure represents in frequency-domain the relative energy of all 9 users (specified by numbers at the right of the sub-figure) after summing the energy over all 100 BS antennas. Given this summation over 100 antennas, the statistics are stable around the expected value of 20 dB and fluctuations are limited to ± 1 dB.

The upper right sub-figure of 3.9 illustrates the frequency-domain correlation of the channel between any two users. There are 45 user pairs for 9 users, identified on the right with the concatenated numbers of the 2 users. For each subcarrier, the correlation is computed as the absolute value of the dot product between the length-100 vectors (corresponding to all BS antennas) of channel coefficients from the two users, divided by the product of the norms of those 2 vectors as normalization. The expected correlation is $1/\sqrt{M} = 0.1$ due to the 100 independent antennas and i.i.d. behavior; the black curve in the middle shows the average over all 45 pairs, which matches this expectation.

Finally, the lower right sub-figure of 3.9 provides 2 additional statistics (shown on the same sub-figure for convenience). The first one, shown between X-axis values 1 and 1200, provides for each subcarrier the condition number (in dB scale) of the product of the (9×100) H matrix by




Figure 3.8: Normalization coefficient spread after averaging over different dimensions (normalized Lund channel).




Figure 3.9: Statistics of a typical Rayleigh channel H (100×9 , 72 i.i.d. taps, frequency analysis over 1200 subcarriers): time-domain delay profile (upper left), frequency-domain response (lower left), cross-correlation between users (upper right), condition number of HH^H (lower right 1–1200) and total energy of the 100 antennas (lower-right 1201–1300).




Figure 3.10: Statistics of a measurement-based Lund channel (100×9) after user normalization.

its conjugate transpose, which is the matrix to invert in order to compute the ZF precoder. The corresponding values are low and stable around 4 dB, which illustrates the benefit of channel hardening, leading well-conditioned $9 \times 9 HH^H$ matrices. The second statistic (between X-axis values 1201 and 1300 and close to value 0) illustrates for each of the 100 antennas the relative channel energy after summing over users and subcarriers. It can be used to check that all antennas play a similar role in the system.

Figure 3.10 shows the same statistics for a Lund channel, after user normalization (for each of the 9 users, the channel energy summed over antennas and subcarriers is made identical to compensate for path loss differences). In the upper-left time-domain representation, the shorter delay spread of the channel is clearly visible, as compared to the Rayleigh case. This translates in smoother channels in frequency domain (lower-left sub-figure), although the total range of variation now spans 5 dB on user 8. An important statistic to understand system performance is given by the upper right sub-figure: the correlation between specific user pairs is much larger than in the case of the Rayleigh model, especially between users 5 and 9 where the correlation is around 0.55 and explains limitations in MRT performance. Finally, the condition number (lower-right) is slightly increased but to 7 dB only, meaning that the ZF implementation is still very stable and benefits from the nice MaMi properties. The antenna energy plot (lower right sub-figure, rightmost part) shows a range of variation of 5 dB instead of 1 dB, between the different antennas. The visible cyclic pattern most likely comes from the use of a circular antenna array configuration in the model (periods corresponding to successive horizontal circles of 32 antennas each in the measurements).

As illustration, Figure 3.11 shows the un-normalized Lund channel. Compared to the normalized case, the time-domain response is similar although the peak-to-average ratio increases. The path loss is best illustrated by the frequency-domain (lower-left) plot where it can be seen




Figure 3.11: Statistics of a measurement-based Lund channel (100×9) without user normalization.

that one user (3) is 10 dB below the others while 2 users (2 and 7) are 10 dB above the others, based on the default parameters and positions in the channel model. Inter-user correlations are identical to the normalized case, given that correlations are normalized by definition. The impact on the condition number is very large (22 dB instead of 7 dB for the normalized case), due to the large amplitude differences between users, which are not compensated for by adding the different antennas. Finally, the range of energy coming from the different antennas now spans 11 dB.

3.4 Conclusions

Simulations have been performed in order to check the performance of the system with respect to three aspects: using a measurement-based channel model versus a theoretical Rayleigh model, using MRT precoding versus ZF precoding, and testing various normalization constraints.

When using Rayleigh channels, all combinations of normalization constraints lead to the same performance except for the most constraining combination of normalization constraints (simultaneously on antennas, users and subcarriers) which leads to a degradation of 1.5 dB in required SNR. Those conclusions are similar for MRT and ZF precoding, with in all cases a 1 dB better performance in the case of ZF.

When using the Lund channel model, the MRT never works properly due to excessive correlation between some users, validating the recommendation that ZF should normally be used except for special cases and theoretical analysis. In absence of normalization constraints, the Lund model gives a significant degradation (close to 10 dB) over the Rayleigh model but this

mostly comes from a difference in path loss (or total channel energy) among the different users. Indeed, the difference in channel energy is 20 dB between the best and the worst user. Problems of different user path loss and power control are crucial in practical deployments but out of the scope of this study. They require specific solutions using link adaptation and scheduling. Moreover, given that the Rayleigh model considers users of equal path loss, a renormalization of the Lund channel was applied such that each user has the same channel energy after summing over antennas and subcarriers, enabling the Rayleigh and Lund models to be compared on an equal path loss basis. This is independent of the precoder normalization options being investigated.

This renormalized Lund channel model gives similar performance to the Rayleigh case when using ZF, i.e., in the selected 100×9 system with QPSK LDPC 3/4, error-free operation around -14 dB SNR. However, more differences were found between normalization options. Especially, when constraining both antennas and users, a large (10 dB) degradation is observed, possibly coming from many weak (antenna, user) combinations in the channel where energy is wasted when normalization is applied, or to a more pronounced suppression of the interference cancellation effect due to normalization on those channels. Other combinations have a limited impact. Hence, as recommendation, the antenna normalization can be used as it corresponds to a physical PA constraint but it should not be combined with user normalization. Alternatively, different geometries than the circular array used in the model can probably improve performance in sectorized operation, e.g., targeting 120 degrees instead of a 360-degrees circular array. The subcarrier normalization can be applied without degradation in order to output a fully frequency-flat spectrum.

This does not mean antenna normalization has to be used: the system generally works best without any normalization constraint but it offers some normalization options for implementation without degrading performance. As a final note, the antenna normalization refers to precoder matrix normalization only. This only leads to normalized expected values for the signals over different antennas. However, actual signals are data dependent and still fluctuate over time. Antenna normalization is hence not sufficient to enable constant-envelope operation, which is a different concept.

Path loss differences between users obviously play a large role in limiting the average system performance. Independently of normalization issues, relevant power control strategies should be further investigated. This goes out of the scope of this study but should be addressed in the future and in combination with specific normalization constraints as well as link adaptation enabling different modulation and coding combinations for the different users.

Chapter 4

Simulation-based performance validation incorporating MAMMOET solutions

This chapter validates from simulations the performance that can be expected from Massive MIMO systems, putting the different components and MAMMOET solutions together. Section 4.1 comments on the achieved performance on the target scenarios of Chapter 2, based on the different elements selected for the final system and scenario. Section 4.2 cross-validates simulation results obtained with the LuMaMi testbed. Finally, Section 4.3 considers the signal quality obtained from the Workpackage 2 modulator and assesses its impact on the overall system performance.

4.1 Overall performance

Figure 4.1 illustrates the BER performance of the baseline scenario proposed in 2.2. It can be related to theoretical predictions in order to quantify the degradation caused by a real environment. For the selected modulation and coding scheme under SISO AWGN channels, a BER of 10^{-5} is achieved at 11.6 dB. In a 100×10 set-up using ZF precoding the array gain allows a reduction of the required SNR to -8 dB, which is confirmed on the Figure when using AWGN channels and perfect CSI. Both 100×10 and 91×1 were selected in order to keep the same number of degrees of freedom ($M - K$) which determine the array gain under ZF precoding.

When using the Massive MIMO channel model and estimated CSI, the system operates around -3 dB. The 5-dB difference mostly comes from two elements. The first one is the actual channel model, based on measurements, causing a 3-dB degradation as compared to the AWGN case. It is mostly related to increased correlation between users. The second one is channel estimation, causing around 2 dB of degradation in this scenario including smoothing (channel interpolation).

While we cannot influence the actual channel, which impacts the performance for both MaMi or other communication techniques, it is important to bound the impact of channel estimation accuracy as it is specific to the MaMi case using channel-based precoders. The 2-dB degradation observed here falls within the range of values observed in Deliverable D3.3, justifying a margin of up to 3 dB for the channel estimation in the overall link budget considered in Section 5.2 in order to scale the output power.




Figure 4.1: MaMi performance for the scenario of Section 2.2 with real MaMi channel model and estimated CSI, compared to scenarios with AWGN channels and/or ideal CSI.

4.2 Comparison to the testbed performance

4.2.1 Testbed scenarios

Several scenarios have been investigated, as described in more detail in MAMMOET deliverable D4.2 [8], using the LuMaMi massive MIMO testbed developed at ULUND. One particular part of the recorded data from these tests has been selected for use together with the MaMi simulator, to verify that channel data measured in tests with the LuMaMi testbed can be used in simulations with the developed MaMi simulator.

The channel data extracted from the LuMaMi testbed are from the *Spatial multiplexing in mobile environments* tests described in Section 3.2 of MAMMOET deliverable D4.2. These tests were performed in Lund together with Bristol University in September of 2016 and were the first MaMi mobility tests ever performed. In particular, we have used data from the *High-mobility – Vehicular and pedestrian speeds UEs* data set, since it represents a challenging scenario previously not studied.

4.2.2 Cross-validation between testbed and simulations

Figure 4.2 illustrates the performance obtained when storing a channel measured in the testbed and plugging it into the simulator. This is different from using the MaMi channel model which is also based on measurements but is still a statistical model building on different measurement scenarios. The testbed channel leads to a performance degradation around 3 dB with respect to the measurement-based channel model, most likely coming from increased correlation between specific users.




Figure 4.2: MaMi performance when using a channel trace from the testbed, as compared to either AWGN or to the MaMi channel model.

It should be noted that the testbed experiment was performed with only 4 independent users, hence the 100×4 configuration on the figure. Other experiments were performed in the testbed with dual-antenna terminals, enabling up to 8 users but having them always paired with 2 antennas close to each other. In such scenarios, the precoder has more difficulties separating pairs of co-located antennas. This leads to an overall degradation of the system performance (SNR requirement) increasing to 7–8 dB, as compared to the channel model. In such scenarios, scheduling such co-located users in different frames would enhance the system performance. Alternatively, if multiple-antenna terminals are used, 2-stream MIMO could be sent towards those terminals. Thanks to channel hardening and to the presence of the MaMi precoder, they would only require a simple (linear) equalizer in order to complement the imperfect stream separation, while simultaneously relaxing the precoder constraints for those 2 streams.

4.3 Simulated performance with hardware non-idealities

In order to assess the impact of the digital modulator (see deliverable D4.3 [9]), the simulator has been adapted by adding at the base station antenna side a random noise generating a distortion comparable to the EVM obtained from the modulator. Given that the achieved EVM is excellent around -45 dBc, no degradation is expected and the simulator was used to also simulate degraded performance modulators in order to identify the bound on EVM such that the performance can be maintained. Figure 4.3 illustrates the impact of different EVM levels. It shows that as soon as the quality of the modulator exceeds -15 to -20 dBc, the performance is equivalent to the ideal case. This illustrates that the Workpackage 2 modulator design is excellent and highly exceeds the minimal requirements for perfect MaMi operation.




Figure 4.3: MaMi performance considering the impact of the digital-RF modulator (-45 dBc EVM) as well as strongly degraded versions of it, as compared to the ideal case.

Chapter 5

Overall power consumption including hardware components

5.1 Power consumption of digital baseband hardware

In order to validate the power consumption of digital baseband processing for MaMi system, we designed and implemented key digital processing blocks (as highlighted in Fig. 5.1) using STMicroelectronics 28nm fully depleted silicon on insulator (FD-SOI) technology. More specifically, we implemented low-latency FFT/IFFT processor to enable fast OFDM (de)modulation, QR-decomposition based ZF precoder to provide accurate down-link beamforming, and Cholesky decomposition based up-link detector to mitigate inter-user interference with wide-range of performance-power trade-off. The detailed algorithm, architecture, and circuit implementation, as well as chip measurement (post-synthesis) results have been presented in Deliverable D3.3 [7]. Here, we would like to summarize (in Table 5.1) the performance of implemented DSP blocks, showing that it is possible to achieve low power digital baseband processing for MaMi system by leveraging the MaMi features and also by conducting algorithm-hardware co-design.

5.2 Overall power consumption

As compared to the power consumption analysis performed in Deliverable D3.2 [5] and [1], the following updates are performed. First a margin on the required output power is added. It is mostly based on the channel estimation penalty. Depending on scenarios, losses between 2 and

Table 5.1: Digital baseband processing blocks implementation result with ST 28nm FD-SOI

Function block	FFT/IFFT processor	ZF precoder	Multi-user detector
Implementation stage	Synthesis	Chip	Chip
Gate count	167kGE	138kGE	148kGE
Clock frequency	500MHz	300MHz	300MHz
Data rate	1GS/s	300Mb/s	300Mb/s
Power consumption	12.3mW	31mW	18mW
Energy efficiency	12.3pJ/sample	60pJ/bit	60pJ/bit




Figure 5.1: Block diagram of OFDM-based massive MIMO digital baseband processing, with implemented function blocks highlighted

4.5 dB were observed after smoothing (Deliverable D3.3), Section 2.1). In the baseline scenario of this deliverable, the degradation was observed to be 2 dB in Section 4.1. An overall 3-dB margin is hence proposed on the system link budget, representative of the channel estimation penalty in MaMi given that a channel-dependent precoder is required. It also includes the minor penalty corresponding to using ZF instead of MRT, as a few degrees of freedom from the antenna array are consumed in order to cancel interference (this degradation is 0.4 dB in 100×10).

Another element concerns channel smoothing/interpolation. In order to operate with the 3-dB margin proposed above, such an advanced channel estimation mechanism is required, hence the related complexity has to be accounted for. It is based on the analysis performed in Deliverable D3.3 and added to the power model.

Table 5.2 summarizes the differences between macro and MaMi base station types. For the rest all common parameters are based on the scenarios in 2.2. The link budget is scaled starting from the 43 dBm per stream value in macro (as it outputs 49 dBm for 4 streams), reducing it based on the array gain from 100 antennas, multiplying by the number of supported users and adding the 3-dB margin discussed above for channel estimation:

$$\begin{aligned} P_{Output,total} &= 43 - 10 \log_{10}(M) + 10 \log_{10}(K) + 3 \\ &= 36 \text{ dBm} \end{aligned} \quad (5.1)$$

$$\begin{aligned} P_{Output,PA} &= P_{Output,total} - 10 \log_{10}(M) \\ &= 16 \text{ dBm} \end{aligned} \quad (5.2)$$

Figure 5.2 illustrates the power consumption estimated from the model for the baseline scenario assuming technology scaling to the year 2016, for macro base stations as well as for 2 types of analog architectures. Three sectors are assumed in both cases. As compared to the study in Deliverable D3.2, feeder losses have also been removed for the reference (macro) base station type, given the remote radio head assumption that can be made in order to remove those losses. This enables a more fair comparison between the different architectures.

As can be seen on the figure, MaMi solutions enable a strong reduction of the power consumption, by a factor 10 to 40 times, depending on the architecture. Simultaneously, in this scenario a slight throughput increase is achieved. Alternatively, when more users are simul-

Table 5.2: Power modeling parameters differing between macro and MaMi base stations. All parameters are given per sector unless specified.

Parameter	Macro	MaMi
Antennas	4	100
Streams	4 (SU-MIMO)	10 (MU-MIMO)
Output per PA	43 dBm	16 dBm
Total output	49 dBm	36 dBm
Quantization	24-bit	4-bit
Throughput (3 sectors)	1.1 Gbps	1.25 Gbps




Figure 5.2: Power model assessment between macro (config. 1), MaMi with traditional hardware (config. 2) and MaMi with advanced digital-RF (config. 3).

Table 5.3: Detailed power consumption for one sector of the MaMi baseline scenario, considering both analog architecture types.

Traditional hardware			
Component	Downlink	Uplink	Training
PA	8.0	0	0
Analog	18.5	27.5	27.5
Digital	0.6	0.6	2.6
Supply	4.9	5.1	5.5
TOTAL	31.4	33.2	35.6
Advanced digital RF			
PA	8.0	0	0
Analog	3.0	3.3	3.3
Digital	0.6	0.6	2.6
Supply	2.1	0.7	1.1
TOTAL	13.1	4.6	7.0

taneously active, MaMi can also provide a further increase in throughput while consequently reducing a bit the savings in power consumption. The detailed power consumption over the different phases and components is given by Table 5.3. The impact of the channel smoothing is visible on the digital part of the training phase but accounts for only 0.5 W for the whole 3-sector base station when averaged over the frame. Figures 5.4 and 5.3 illustrate graphically this split over the different phases and the relative share of the different components in the total power consumption.

Let us consider the analog modulator consumption as compared to the Workpackage 2 design. By taking optimistic assumptions on how far the digital-RF approach could be applied and optimized, the power model assumes that analog power for 100 antennas could be reduced from 18.5 W to 3 W, i.e., 30 mW per antenna chain. Considering the real hardware, the measured power consumption of the Workpackage 2 modulator design is around 100 mW (Deliverable D4.3 [9]). Adding the required DSP pre-processing such as RF upsampling could increase the total power consumption to around 250 mW per antenna chain. Those values are closer to the traditional architectures in the power model than the digital-RF projections used, but those projections may be a bit too optimistic. More importantly, the offered resolution from the hardware modulator was found to be much higher than required. It offers a very high accuracy with an EVM around -45 dBc and could be used in much broader applications than Massive MIMO, where -15 or -20 dBc could be sufficient. Given the very relaxed specs required in Massive MIMO, a specific redesign could strongly reduce the quantization and modulator resolution, leading to substantial additional power savings and explaining the difference. This is at least true for data channel considered all over this project. The design and optimization of a MaMi system also covering broadcasting and control channels is another challenging topic which should be investigated by future projects as those channels have different requirements in order to maintain the cell coverage.




Figure 5.3: Power split over the 3 phases when using traditional hardware.




Figure 5.4: Power split over the 3 phases when using advanced digital RF.

Chapter 6

Conclusions

Over the three years of the MAMMOET project, Massive MIMO has been brought from a concept to a working solution for cellular networks, offering higher capacity, simpler components and reduced power consumption. A number of elements were only possible thanks to a broad multidisciplinary approach, combining theory, measurements, modeling, simulations, testbeds and hardware design.

MaMi was shown to be particularly suited to 5G scenarios such as open exhibition or enhanced Mobile BroadBand, where it can advantageously replace traditional base stations and be dimensioned to offer benefits in capacity and/or power consumption.

New channel models generated from MAMMOET measurements play a crucial role in the investigation of the system performance. Compared to more traditional theoretical model, they can have a large impact on the simulated performance. Differences in path loss between users require a specific focus, in line with power control and scheduling strategies. Normalization constraints also have a specific impact on performance. When those effects are properly addressed, an excellent performance is obtained for the target scenarios.

The main limitations come from co-located users, having a very strong correlation between their respective channels, and channel estimation as proper channel knowledge is used to compute the MaMi precoder. The channel estimation noise can however be reduced with proper smoothing strategies investigated in MAMMOET.

The hardware modulator design was found to yield excellent performance. Thanks to the robustness of MaMi communication, the signal quality generated by the modulator was found to be much better than the minimal requirements, opening the door to further reductions in hardware complexity and power consumption.

Considering the complete base station, the power reduction promise from MaMi was assessed through high-level modeling and confirmed from designs of both RF modulator and digital base-band components, leading an overall 10-fold reduction or more on the total power consumption for the target scenarios.

Simulation-based performance assessment was performed to include actual algorithmic solutions and measurement results for channels and innovative hardware components. We can conclude that the technological progress achieved is considerable, and the confidence level for Massive MIMO has been raised significantly, both with respect to its capability to increase capacity (spectral efficiency) and to operate at superior energy efficiency.

Chapter 7

Abbreviations

16-QAM	16-point Quadrature Amplitude Modulation
3D	3-Dimension
3GPP	3rd Generation Partnership Project
5G	5th Generation
A	Antenna (normalization)
AWGN	Additive White Gaussian Noise
BER	Bit Error Rate
BS	Base Station
CP	Cyclic Prefix
CSI	Channel State Information
DSP	Digital Signal Processing
eMBB	enhanced Mobile BroadBand
EVM	Error Vector Magnitude
FD-SOI	Fully-Depleted Silicon On Insulator
FFT	Fast Fourier Transform
IFFT	Inverse Fast Fourier Transform
LDPC	Low-Density Parity-check Code
LTE	Long-Term Evolution
MaMi	Massive MIMO
MIMO	Multiple-Input Multiple-Output
MMSE	Minimum Mean-Squared Error
MRC	Maximum-Ratio Combining
MRT	Maximum-Ratio Transmission
MU	Multi-User
NLOS	Non Line-Of-Sight
NR	New Radio of 5G
OFDM	Orthogonal Frequency-Division Multiplexing
PA	Power Amplifier

QPSK	Quadrature Phase Shift Keying
RF	Radio Frequency
S	Subcarrier (normalization)
SINR	Signal to Interference plus Noise Ratio
SISO	Single-Input Single-Output
SNR	Signal-to-Noise Ratio
SU	Single-User
TDD	Time-Domain Duplexing
U	User (normalization)
ZF	Zero-Forcing

Bibliography

- [1] Claude Desset and Björn Debaillie. Massive MIMO for energy-efficient communications. In *EuMC*, London, UK, October 2016.
- [2] FP7 MAMMOET. *Deliverable D1.1, System scenarios and requirements specifications*. December 2014.
- [3] FP7 MAMMOET. *Deliverable D1.2, MaMi Channel Characteristics: Measurement Results*. June 2015.
- [4] FP7 MAMMOET. *Deliverable D1.3, MaMi Channel Characteristics: MAMMOET Channel*. June 2015.
- [5] FP7 MAMMOET. *Deliverable D3.2, Distributed and centralized baseband processing algorithms, architectures, and platforms*. December 2015.
- [6] FP7 MAMMOET. *Deliverable D4.1, Validation simulation framework*. December 2015.
- [7] FP7 MAMMOET. *Deliverable D3.3, Hardware-aware signal processing for MaMi systems*. December 2016.
- [8] FP7 MAMMOET. *Deliverable D4.2, Test-bed based assessment and proof of concept*. September 2016.
- [9] FP7 MAMMOET. *Deliverable D4.3, Proof of Concept of Innovative Hardware Including Transmitter Characterisation*. December 2016.