

TECHNISCHE UNIVERSITÄT MÜNCHEN
Lehrstuhl für Kommunikationsnetze

Radio Resource Management and System Spectral Efficiency in LTE Multi-Cell Systems

Dipl.-Inform. Jan Karl Ellenbeck, MBA

Vollständiger Abdruck der von der Fakultät für Elektrotechnik und Informationstechnik
der Technischen Universität München zur Erlangung des akademischen Grades eines

Doktor-Ingenieurs (Dr.-Ing.)

genehmigten Dissertation.

Vorsitzender: Univ.-Prof. Dr. sc. techn. Gerhard Kramer
Prüfer der Dissertation: 1. Univ.-Prof. Dr.-Ing. Jörg Eberspächer (i.R.)
2. Univ.-Prof. Dr.-Ing. Wolfgang Utschick

Die Dissertation wurde am 30.04.2014 bei der Technischen Universität München ein-
gereicht und durch die Fakultät für Elektrotechnik und Informationstechnik am 24.02.2015
angenommen.

TO MY FAMILY

Acknowledgements

When I started my work as a research assistant at the Institute of Communication Networks under the supervision of Professor Dr. Jörg Eberspächer in December 2006, the first release of the LTE standard was still two years away. At the time the thesis was submitted, four releases of the standard had been completed and LTE networks were deployed in more than 100 countries covering 20 % of the world population. In light of this rapid progress in both the industry and the research community, developing simulation tools, conducting research, and, finally, completing this thesis would not have been possible without the help of numerous people.

First and most importantly, I am sincerely grateful to my advisor Professor Eberspächer for his constant trust, support, and for giving me the freedom research needs to thrive. He created a cooperative spirit that allowed for fruitful exchanges between colleagues working in different research fields. Highlights were also the regular institute events. The annual weekend seminars, for example, brought together the current research assistants and the "oldies", who shared their experiences and perspectives from their careers after leaving the institute.

I would also like to thank Professor Dr. Wolfgang Utschick for serving as the second examiner. I learned a lot about optimization and multi-antenna systems from attending the lectures and talks at his institute and through the collaboration with Andreas Dotzler from his group. Many thanks also go to Professor Dr. Gerhard Kramer for chairing the examining committee.

Among the colleagues at the institute, my first thanks go to the members of the mobile communications group. I am indebted to Dr. habil. Christian Hartmann, who provided guidance from the very beginning, and to Dr. Ulrike Buhl for the many helpful discussions and the careful review of the manuscript. I would also like to thank Dr. Moritz Kiese, Dr. Carmen Mas Machuca, Dr. Silke Meister, Dr. Robert Prinz, Dr. Matthias Scheffel, Dr. Robert Vilzmann, and Dr. Hans-Martin Zimmermann for sharing their knowledge and opinions during the many fruitful technical and non-technical discussions we had.

As the key tool for the research in this thesis, I developed the *IMTAphy* LTE system-level simulator. I want to thank Dr. Christoph Spleiß and Dr. Oliver Hanka for their help with the in-house simulation cluster without which it would have been impossible to perform the vast amount of simulations for this thesis. Not only in this context, I could rely on the support of Dr. Maier, who was always able to provide the authoritative answer to administrative questions of any kind. *IMTAphy* is based on core libraries and frameworks from the openWNS simulator developed by the Communication Networks (ComNets) research group at RWTH Aachen, where I wrote my Diploma thesis before coming to Munich. My thanks go to Daniel Bültmann and Maciej Mühleisen for their support on openWNS issues and to Professor Dr. Bernhard Walke for releasing the openWNS as open-source software upon which I could build *IMTAphy*.

I had the pleasure to supervise a number of students who contributed to my research and to the development of the *IMTAphy* simulator and its predecessors as part of their

final thesis or as student workers. My appreciation goes to Zhong Lei, Rui Hu, Dr. Hussein Al Shatri, Fernando Andrés Quiroga, Boyan Lazarov, Mahmoud Hammoud, Rocío Casas Villagarcia, Johannes Schmidt, Obaid Mushtaq, Farid Sheikh, Safiullah Shahrukh Qazi, Dincer Beken, and Dmytro Bobkov. Since `IMTAphy` has been made open-source, our work lives on and I hope the simulator and this thesis will be helpful to other researchers of 4G and also 5G systems.

During my work at Intel, I got to know the performance of LTE systems from a very practical point of view when helping to optimize the lab and field performance of Intel's latest cellular modems. In this regard, I am thankful to Dr. Christian Drewes, Dr. Thorsten Clevorn, and Professor Dr. Josef Hausner for sharing their long experience and for providing such a great work environment. I also benefited greatly from Sabine Rössel's profound insight into the LTE standard.

Last but not least, I would like to thank my family and friends for their support and encouragement along the way. Above all, I want to thank Monika for her endless patience and unwavering support, especially in the last stretches of writing this thesis.

Grünwald, August 2015

Jan Ellenbeck

Abstract

In this thesis, the system spectral efficiency of LTE multi-cell networks is evaluated. To this end, an efficient system-level simulator is presented, whose radio channel modeling and link-to-system mapping approach are discussed in detail. The holistic analysis of different LTE transmission modes and radio resource management (RRM) schemes reveals dependencies between RRM schemes and scenario conditions. Finally, the performance of selected interference management and multi-user multiple-input multiple-output (MIMO) schemes is evaluated.

Kurzfassung

Das Ziel der Arbeit ist die Leistungsbewertung von zellularen LTE Mobilfunksystemen hinsichtlich ihrer spektralen Effizienz. Der hierzu entwickelte effiziente System-Level-Simulator und insbesondere die Modellierung des Funkkanals und der physikalischen Schicht werden im Detail vorgestellt. Die verschiedenen Übertragungsmodi werden ganzheitlich analysiert mit einem Hauptaugenmerk auf dem Zusammenspiel der Verfahren zur Funkressourcenverwaltung und der Szenarioeinflüsse. Schließlich werden ausgewählte Interferenzmanagement- und Mehrnutzer Multiple-Input Multiple-Output (MIMO)-Verfahren untersucht.

Contents

1. Introduction	1
1.1. Motivation	1
1.2. Contributions	6
1.3. Outline of the Thesis	7
2. From Channel Capacities to Evaluating System Spectral Efficiency	9
2.1. An Introduction to Theoretical Channel Capacities	9
2.1.1. System Model for a Single Wireless MIMO Link	10
2.1.2. Single Link MIMO Channel Capacity	14
2.1.3. From Single Links to Multi-Cell Systems	18
2.2. LTE System Design	23
2.2.1. Overview of the LTE Protocol Stack	24
2.2.2. The LTE Physical Layer	28
2.2.3. Radio Resource Management in LTE	33
2.3. LTE Link Spectral Efficiency	41
2.3.1. Suboptimal Transmitters and Receivers	42
2.3.2. Overhead in the LTE Physical Resource Grid	44
2.3.3. Total PHY Loss and Link Spectral Efficiency	45
2.4. Evaluating System Spectral Efficiency by System-Level Simulation	46
2.4.1. System-Level RRM Spectral Efficiency Determinants in LTE	47
2.4.2. System-Level Evaluation Methodology	48
2.5. System-Level Simulation Tools	51
2.5.1. A Short Survey of Simulation Tools	51
2.5.2. Overview of IMT _A phy	55
2.6. Summary and Contributions	58
3. The IMT-Advanced Channel Models	59
3.1. Test Environments and Deployment Scenarios	59
3.2. Antenna Modeling	62
3.2.1. IMT-Advanced Antenna Field Pattern Model	63
3.2.2. 3GPP Antenna Array Layout and Polarization Model	64
3.3. The IMT-Advanced Channel Models	66
3.3.1. Large-Scale Link Properties	68
3.3.2. Initialization of Small-Scale Parameters	72
3.3.3. Coefficient Generation	73
3.3.4. Computationally Efficient Time Evolution of CIRs and CTFs	75
3.4. Channel Model Calibration	76
3.4.1. Large-Scale Calibration Metrics	77

3.4.2.	Small-Scale Calibration Metrics	79
3.4.3.	CIR and CTF Calibrations	80
3.5.	Summary and Contributions	81
4.	Link-to-System Modeling for LTE System-Level Simulations	83
4.1.	Overview of Link-to-System Modeling for System-Level Simulations . .	83
4.1.1.	Link-Level Simulations	83
4.1.2.	Link Performance Models in the Literature	84
4.1.3.	The Standard Link-to-System Modeling Concept	86
4.2.	Modeling of MIMO Linear Receivers and Precoders	89
4.2.1.	SINR Computation Per MIMO Layer	90
4.2.2.	Linear MIMO Receive Filters	91
4.2.3.	Transmit Diversity using Space-Frequency Block Codes	92
4.3.	Modeling Channel Estimation Errors	95
4.3.1.	Modeling Serving Channel Matrix Estimation Errors	96
4.3.2.	Modeling IPN Covariance Matrix Estimation Errors	97
4.4.	Link Performance Model	99
4.4.1.	SINR Compression to Compute Effective SINR Values	99
4.4.2.	Block Error Modeling	101
4.5.	3GPP IMT-A System-Level Simulator Calibration	103
4.5.1.	Downlink Simulation Assumptions	103
4.5.2.	Uplink Simulation Assumptions	104
4.5.3.	Simulator Calibration Results	105
4.6.	Summary and Contributions	108
5.	Performance of Fundamental RRM Schemes	109
5.1.	Scheduling and Link Adaptation (SISO)	109
5.1.1.	Channel-Adaptive Proportional Fair Downlink Scheduling	110
5.1.2.	Challenges of Outdated CSI	118
5.2.	Spatial Processing at the Receiver (SIMO)	125
5.2.1.	Performance Gains from Multiple Rx Antennas	126
5.2.2.	Serving Channel and IPN Estimation Impairments	129
5.3.	SFBC Transmit Diversity in LTE	133
5.3.1.	Higher Robustness for Unscheduled Low-Rate Transmissions . .	134
5.3.2.	Disadvantages of SFBC with Channel-Adaptive Scheduling . . .	135
5.4.	Beamforming at the BS (MISO)	137
5.4.1.	Codebook-Based Beamforming in LTE (TM6)	138
5.4.2.	Performance of Codebook-Based Beamforming	142
5.4.3.	Link Adaptation for Beamforming in Multi-Cell Systems	146
5.5.	Spatial Multiplexing (MIMO)	148
5.5.1.	Open-Loop and Closed-Loop Spatial Multiplexing in LTE	148
5.5.2.	Performance of Spatial Multiplexing with TM3 and TM4 in LTE .	151
5.6.	Comparing LTE Transmission Modes 1, 2, 3, 4, and 6	158
5.7.	Area Spectral Efficiency	164
5.8.	Summary and Contributions	167

6. Selected Single- and Multi-Cell RRM Methods	169
6.1. Frequency Reuse Based Interference Management	170
6.1.1. Introduction to Frequency Reuse Schemes	170
6.1.2. Performance Evaluation of Fractional Frequency Reuse in LTE . .	173
6.1.3. Dynamic Frequency Reuse Coordination — An Excursion	179
6.2. Neighbor-Considerate and Multi-User MIMO Transmission Schemes . .	187
6.2.1. Predictable Beamforming	187
6.2.2. Multi-User MIMO: Zero-Forcing and PU ² RC	189
6.2.3. Performance Evaluation of ZF and PU ² RC MU-MIMO	194
6.3. Outlook to Future Improvements in LTE-Advanced and 5G Systems . .	202
6.3.1. Deployment and Spectrum Aspects	202
6.3.2. Link Efficiency Aspects	203
6.3.3. System-Level RRM Efficiency Aspects	204
6.4. Summary and Contributions	208
7. Conclusions	209
A. Additional Simulation Results for Chapter 5	213
B. Codebooks from the LTE Standard	221
C. Acronyms	229
Nomenclature	235
List of Figures	243
List of Tables	249
Publications and Presentations by the Author	251
General Literature	253
Standards and Standardization Documents	269
Cited Websites	275

1. Introduction

1.1. Motivation

In recent years, there has been a tremendous growth in the usage of cellular mobile networks. According to estimates by the International Telecommunication Union (ITU), in 2013 the global penetration of mobile cellular subscriptions has reached 6.8 billion or 96 % of the global population [Int13]. At the same time, the number of mobile broadband data subscriptions has grown at an annual rate of 40 % since 2007 to reach 2.1 billion subscriptions at the end of 2013. Together with the number of broadband subscriptions, also the demand for data rates can be expected to increase further in the future as smartphones and the mobile consumption of multimedia content become more and more popular. For example, in its global traffic forecast [Cis14], Cisco estimates the worldwide mobile data traffic to increase 11-fold from 2013 to 2018.

The availability of radio spectrum that is physically suitable for cellular mobile communications and not yet allocated otherwise is limited. This is despite recent efforts by the 2012 World Radiocommunication Conference (WRC), which, for example, plans to make the 700 MHz band available for worldwide mobile communications use by 2015. Due to the scarcity of the radio spectral resource, the operators of cellular networks often have to pay spectrum licensing costs in the billions of Euros. Consequentially, operators are interested in deploying radio access technologies that make the most efficient use of the spectrum.

Evolution of Cellular Standards towards LTE and LTE-Advanced

In order to cope with the growing demand for data rates in mobile networks, the standardization organizations, like the 3rd Generation Partnership Project (3GPP), have continuously improved cellular standards and introduced new ones. Figure 1.1 shows an overview of 3GPP standards and their peak data rates. With every new technology, higher peak rates can be achieved. The first digital systems, known as the second generation (2G) with the Global System for Mobile Communication (GSM) being the most prominent member, were introduced in the early 1990s [EVBH08]. The General Packet Radio Service (GPRS) extension to GSM allowed the provision of packet-switched data service but with a modest peak rate as shown in Fig. 1.1. To promote higher rates and efficient spectrum use, the ITU initiated the International Mobile Telecommunications (IMT) framework with the IMT-2000 system, which encompasses all technologies that are officially considered to be the third generation (3G). Among them is the Enhanced Data Rates for GSM Evolution (EDGE) extension of GSM. But the most

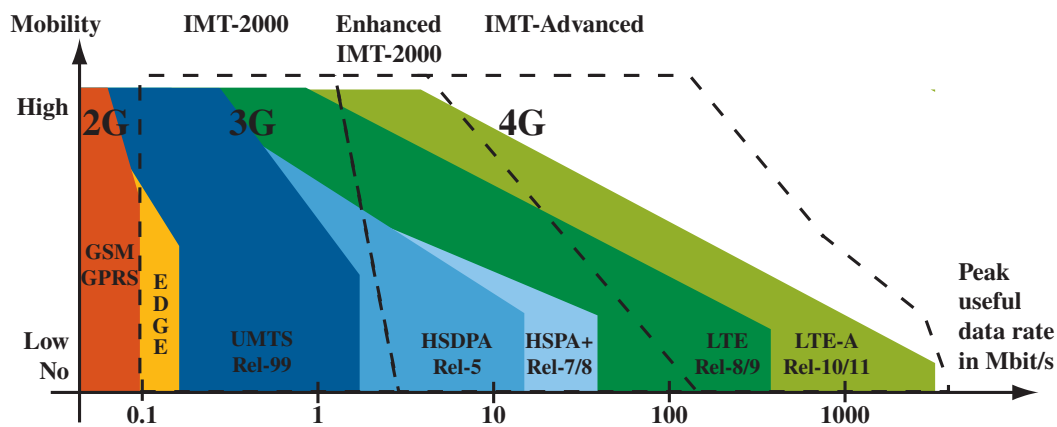


Figure 1.1.: Overview of 3GPP cellular standards and their peak data rates in the context of the IMT-2000 and IMT-Advanced classification [M.1645; Sch08]

prominent 3G standards are the Universal Mobile Telecommunication System (UMTS) standard and its High Speed Packet Access (HSPA) evolutions that aim at further data rate increases. The most advanced HSPA+ versions, as well as the Long Term Evolution (LTE) standard in the releases 8 and 9, are officially considered to be enhanced IMT-2000 and thus 3G standards. However, the ITU acknowledged the market realities [Int10b] and allows already those standards to be considered as 4G. Originally, the 4G designation was reserved for systems that fulfill the ITU's IMT-Advanced requirements [M.2134], which were established in 2008.

With the advent of IMT-Advanced, the number of mobile standards has converged. As of 2014, only two contenders have been approved as IMT-Advanced technologies [Int10a]: the IEEE 802.16m standard (not shown in Fig. 1.1) and the LTE-Advanced standard, which stands for the Release-10 and beyond of the LTE standard developed by the 3GPP. While the race between 802.16m and LTE seemed to be open at first, the momentum has decidedly shifted in LTE's favor. In this thesis, we will thus focus on the initial LTE Release-8 and on advanced schemes possible with later LTE-Advanced releases. Due to its technological similarity, many results also apply to the 802.16m standard.

For a number of reasons, Fig. 1.1 serves as a motivation for the approach followed in this thesis. First, the figure shows peak rates, which usually cannot be achieved in practice. It is thus of interest to evaluate a rate that can be achieved on average under representative assumptions. Second, the impression of a 10 000-fold gain between GPRS and LTE-Advanced is misleading because for the shown peak rate, the 3GPP assumes a 500-fold frequency bandwidth (200 kHz versus 100 MHz). This motivates the use of rates normalized to the employed spectrum for a fair comparison. Third, we observe a trade-off between the achievable rates and the user mobility scenario. As we will see, a system can adapt to environmental factors, like the user mobility, by means of so-called *radio resource management* (RRM) schemes. Advanced RRM schemes, as we will discuss in this thesis, provide the levers to increase the data rates beyond the bandwidth-dependent scaling.




Figure 1.2.: Exemplary cellular system with 7 base station sites serving 21 cell sectors

Spectral Efficiency of Cellular Systems

The central figure of merit for cellular operators, which is in the focus of this thesis, is thus the so-called *system spectral efficiency* (or also *cell spectral efficiency*) η . It indicates how much traffic a given radio access technology can serve on average to the users in one cell of the cellular system. The spectral efficiency measures the correctly delivered bits per second and per utilized Hertz of frequency spectrum [M.2134]:

$$\eta = \frac{\text{correctly delivered bits system-wide}}{\text{system bandwidth} \cdot \text{time} \cdot \text{number of cells}} \text{ in bit/s/Hz/cell.} \quad (1.1)$$

We distinguish between the spectral efficiency in the *downlink* (from the network towards the user) and *uplink* direction (from the user towards the network). In this thesis, we will focus on the downlink because the capacity demand is usually higher in the downlink than in the uplink.

In a cellular system, as depicted in Fig. 1.2, a number of *base stations* provide the users in the coverage area of the system with wireless access to send and receive data to and from some kind of backbone network like, e.g., the Internet. The users of the cellular system, which we will call *mobile stations* throughout this thesis, can be located at arbitrary positions in the coverage area. A mobile station (MS), depicted as a handset in Fig. 1.2, is usually only connected to one base station (BS). The part of the coverage area containing all mobiles served by the same base station is called a cell. Following the tradition in modeling cellular networks, we draw the cells as hexagons and place the base station sites at regular positions in the scenario. We note that the exact shape of the cells depends on signal propagation conditions. In real-world deployments, base station sites have to be chosen according to terrain and urban conditions. Figure 1.2 shows a sectorized deployment where from the site of each base station tower three cell sectors are served by a linear array of four antennas per sector.

The spectral efficiency defined in (1.1) is the aggregate efficiency over all BS-MS links in the system. It is normalized to the system's frequency bandwidth because the throughput scales linearly with the available spectrum. The maximum normalized rate per




Figure 1.3.: Cell spectral efficiency distribution example

link, called the *peak spectral efficiency*, can usually be reached only under idealistic conditions and when all resources are assigned to one user. In a realistic setting though, the resources have to be shared by the users of one cell and a user's throughput depends on its individual radio conditions. It is thus common to evaluate the system-wide distribution of the *normalized user throughput* γ_i defined as [M.2134]:

$$\gamma_i = \frac{\text{correctly delivered bits for user } i}{\text{system bandwidth} \cdot \text{time}} \text{ in bit/s/Hz.} \quad (1.2)$$

Figure 1.3 shows exemplary cumulative distribution functions (CDFs) of the system-wide normalized user throughput distribution and of the total cell throughput. The former depends on the actual numbers of users per cell. Based on these two spectral efficiency distributions, the two most important system-level key performance indicators (KPIs) are defined: the *system spectral efficiency* introduced in (1.1) is the mean of the total normalized cell throughput, and the *cell-edge user spectral efficiency* (or *cell-edge user throughput*) is defined as the 5% level of the system-wide per-user normalized throughput distribution. Having this dual objective assures that the system does not increase the system spectral efficiency in an unfair way by only serving those users that can achieve the highest data rates while starving the worst-off users.

Together with the requirements [M.2134] for IMT-Advanced systems, the ITU also defined the evaluation guidelines for validating candidate technologies. Most notably, the evaluation guidelines in report [M.2135] define a sophisticated spatial channel model and require the use of system-level simulations to assess the system spectral efficiency. Due to the authority of the ITU, these evaluation requirements quickly became a de-facto standard for the performance evaluation of cellular systems. Consequently, almost all quantitative results presented in this thesis were obtained by system-level simulations according to the ITU evaluation guidelines.

Radio Resource Management

In order to achieve a high spectral efficiency, a wide range of RRM methods are used in cellular systems. They control and adapt the usage of spectral resources depending




Figure 1.4.: High-level overview of determinants affecting system-level data rates

on the needs of the users and the constraints of the deployment and the radio access technology. Within one cell, the BS employs RRM schemes in order to distribute the available radio resources among its users, to transmit to each user in a way that is suitable for the radio propagation conditions on each link, and to make best use of the technical equipment, like antenna arrays. Between the cells of a cellular system, RRM schemes aim at controlling the so-called *inter-cell interference*. Another inter-cell aspect of RRM, which is not in the scope of this thesis, is providing a handover mechanism that maintains the connection to a user while the user travels across the system from one cell to another. Being in control of the spectral resources in the system, the various RRM schemes and their interactions determine the spectral efficiency of the system and also the trade-offs between, for example, the cell-edge user throughput and the overall system spectral efficiency. The protocols and interfaces between components of a radio access technology are usually standardized to allow interoperability between different vendors. But many algorithmic details and parameters of the RRM schemes are not standardized. They rather offer a way for each vendor to differentiate from the competition. Analyzing the dependencies between RRM methods and their contributions to the spectral efficiency KPIs will be the central focus of this thesis.

System-Level Viewpoint

Figure 1.4 gives a high-level overview of the determinants affecting the available data rates on a system level. In later chapters, we will expand the shaded boxes and take a closer look at the individual factors. The *link spectral efficiency* and what we call the *system-level RRM efficiency*, printed in **bold**, are considered to be under the control of a system engineer. We will thus evaluate which design choices lead to high data rates. The remaining factors, printed in *italics*, we consider as given. A system engineer cannot change them but has to consider their repercussions on the system so that careful modeling is needed. Finally, the unshaded boxes are not considered further: The amount of available spectrum is not important for our evaluation since we will focus on the spectral efficiency as defined in (1.1) in the following. The number of sites, which determines the cell sizes if the system area is fixed, results in a measure called *area spectral efficiency*. In Section 5.7 we will briefly discuss why the cell spectral efficiency is roughly independent of the cell size, which allows us to disregard the cell size and stick to scenario-specific default cell sizes in the remainder of the thesis.

1.2. Contributions

This thesis provides contributions in the following areas:

- **An efficient and calibrated simulation tool:** For the simulation-based performance evaluation of LTE RRM methods, we present the necessary deployment scenarios, spatial channel models, a link-to-system model, as well as models for the lower layers of the LTE user plane. All of the above, and all RRM methods covered in this thesis, are implemented in a C++ system-level simulator called `IMTApHy`. The tool has been extensively calibrated against reference results from industry-leading 3GPP member companies. It is provided as open-source software, allowing interested researchers to reproduce and extend all of our results.
- **Spectral efficiency evaluation of fundamental RRM methods in LTE:** This thesis offers a holistic and consistent evaluation of the interplay of different RRM methods and system parameters in an LTE system. Most results in the literature focus on a single kind of RRM mechanism evaluated in some often arbitrary setup. We argue that most RRM methods cannot be judged in isolation but have to be seen in context with system parameters and other RRM methods they have to interwork with. In particular, we highlight the dependency of the system performance on the structure of the inter-cell interference.
- **Fractional frequency reuse in combination with LTE RRM schemes:** We evaluate the performance of static interference management schemes based on a reduced frequency reuse in the IMT-Advanced scenarios. We show that fractional frequency reuse (FFR), which received considerable attention in the literature, cannot deliver the expected system spectral efficiency or even cell-edge throughput gains if standard LTE RRM schemes are also taken into account. This underlines the importance of adequate modeling and a holistic evaluation.
- **Autonomous interference coordination for femtocell deployments:** In unstructured femtocell deployments, in contrast, dynamic inter-cell frequency reuse coordination can be beneficial, but a central coordination entity is often missing. To this end, we show how autonomous frequency reuse decisions by selfish cells can lead to a suboptimal but stable frequency reuse pattern.
- **Predictable beamforming:** If a base station uses antenna arrays for beamforming in the downlink, the performance of neighboring cells can suffer due to the resulting unpredictable interference pattern. By making the beamforming predictable, we can show that despite a loss in diversity, the overall system performance increases due to better link-adaptation and an implicit coordination gain.
- **MU-MIMO schemes with limited-feedback:** Single- and multi-user multiple-input multiple-output (MIMO) schemes are often evaluated in single-cell settings only. In contrast, we provide a full system-level performance evaluation of two advanced multi-user MIMO (MU-MIMO) schemes that rely on limited feedback comparable to that of standard LTE MIMO schemes. We focus on system aspects such as link adaptation and inter-cell interference and show that MU-MIMO can provide a significant increase in system and cell-edge spectral efficiency.

A part of the results has previously been published in the proceedings of international conferences [Ell07a; Ell07b; Ell08; Ell09a; Ell09b; Ell10; Ell11; Ell12b] and was presented at workshops [EllPres09; EllPres10a; EllPres10b; EllPres11]. In addition, a significant part of the material presented in Chapters 3 and 4 was included in a book chapter [Ell12a] published by John Wiley & Sons.

1.3. Outline of the Thesis

The thesis is divided into the three main areas shown in Fig. 1.5:

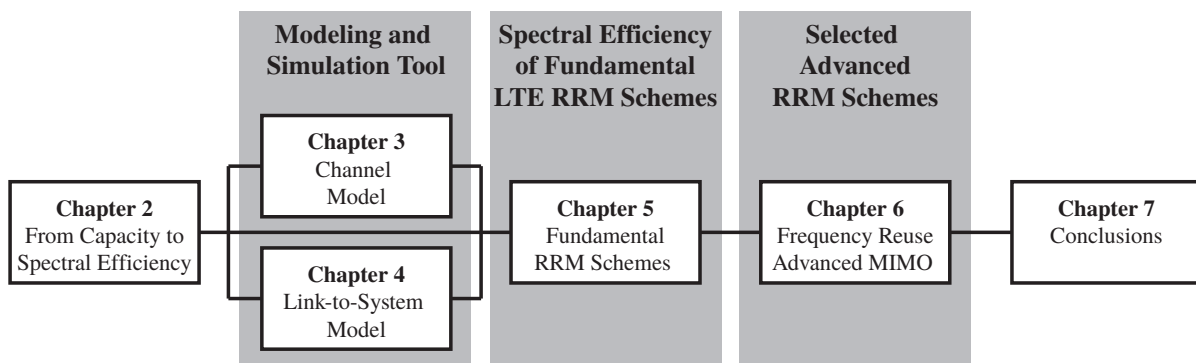


Figure 1.5.: The chapters and the three pillars of the thesis

Chapter 2 lays the foundations for the remainder of the thesis. It introduces fundamental concepts like the single-link channel capacity, the cellular principle, and MIMO schemes. We highlight the contrast between classical information-theoretic approaches, which focus on capacity results, and the notion of system spectral efficiency, which is in the focus of this thesis: actual systems have to deal with practical limitations like limited processing power, delays, measurement inaccuracies, signaling constraints, the mobility and unpredictable behavior of their many users, etc. Thus, their performance cannot be analyzed solely by information-theoretic means. Instead, system performance should be evaluated by means of computer simulation for which the IMT-Advanced guidelines serve as a de-facto standard. In this chapter, we also present some of the design choices for the lower layers of LTE. We give an overview of fundamental LTE RRM mechanisms and discuss in detail the physical layer overhead and other factors that limit the achievable data rates in LTE systems. Finally, we briefly review some of the system-level simulation tools that are available to academia. Their limitations motivate the development of our own simulation environment, whose underlying models are detailed in the following two chapters.

Chapter 3: Modern cellular systems like LTE do not solely aim at mitigating the fading characteristics of the typical multi-path cellular propagation channels but aim at exploiting the characteristics of the propagation channels by means of MIMO and channel-adaptive transmission. Consequently, a detailed and realistic scenario and channel model is a basic requirement for the performance evaluation of LTE. In this

chapter, we present the IMT-Advanced deployment scenarios and the stochastic geometry-based channel model with its scenario-specific parameters. We discuss how to efficiently implement the channel generation in a simulation tool. We also present the various calibration mechanisms that we used to validate our implementation against results provided by IMT-Advanced evaluation groups and 3GPP members.

In **Chapter 4** we present so-called link-to-system or PHY abstraction methods that allow us to model the performance of a link in a system-level simulation without the overhead that a detailed link-level simulation would entail. We use the signal to interference plus noise ratio (SINR) as the central figure of merit to derive block error ratios with the help of pre-computed mapping curves. We discuss how to model SINRs if MIMO reception with linear receivers, precoding at the transmitter, or an Alamouti transmit diversity scheme are used. In addition, we discuss how channel estimation errors impairing the performance in an actual system can be modeled in this context. The link-to-system model is implemented in our simulation tool. We underline the correct implementation and the ability to generate meaningful results by showing the calibration of the link-to-system model against reference calibration data from industry-leading 3GPP partner companies.

Chapter 5 introduces basic LTE RRM methods like feedback-based link adaptation, outer-loop link adaptation, channel adaptive scheduling, transmitter-side spatial processing, spatial processing at the receiver, and open-loop and closed-loop single-user spatial multiplexing MIMO operation. Based on suitable combinations of these methods, we compare the performance of the most important LTE transmission modes. One key contribution and insight discussed in this chapter is the high interdependency of RRM methods and system design parameters that make it impossible to evaluate the performance of an RRM method on its own. The analysis of the presented basic RRM schemes suggests possible RRM enhancements. Two promising directions are further pursued in the following chapter.

In the first part of **Chapter 6**, we consider interference management schemes that reduce the inter-cell interference, which fundamentally limits the spectral efficiency of every cellular system, by a coordinated reduced frequency reuse across neighboring cells. We show that a fractional reuse scheme is inferior to a full reuse if channel-adaptive scheduling and multiple antennas are used in the system. As part of a detour to femtocell deployments, we briefly analyze a dynamic reuse coordination scheme that achieves a system-wide interference reduction by autonomous reuse decisions. In the second part of the chapter, we consider a more promising approach to increasing spectral efficiency than a reduced reuse. First, we examine what we call a neighbor-considerate beamforming method that helps the link adaptation by making the inter-cell interference more predictable. This is also an aspect of the MU-MIMO schemes we evaluate in our last contribution. We show that with only limited channel state feedback, considerable spectral efficiency gains can be achieved over single-user MIMO.

Finally, **Chapter 7** summarizes the thesis and its main contributions.

An overview of the mathematical notation and all symbols used in this thesis can be found in the Appendix on page 235.

2. From Channel Capacities to Evaluating System Spectral Efficiency

In this chapter, we introduce the basic methods required for the evaluation of spectral efficiency and radio resource management schemes in LTE systems. First, we give an overview of classical wireless communication concepts and present some theoretical capacity results. Then we introduce the most important aspects of the LTE physical layer and protocol stack and briefly introduce possible RRM schemes for LTE. Based on that, we discuss reasons why practical systems do not achieve theoretical channel capacities and motivate the need for a system-level simulation approach as it is applied in this thesis. Finally, we give a short overview of existing simulation tools and the IMT_Aphy tool developed for this thesis.

While some of the general concepts presented in this chapter apply to both the uplink and downlink direction, we concentrate on the downlink. Most of the discussed features were introduced in the initial release of the LTE standard, while some of the RRM schemes are only available since the LTE-Advanced releases of the standard.

2.1. An Introduction to Theoretical Channel Capacities

In the introduction of this thesis, we already defined the *system spectral efficiency* in equation (1.1). As a measure of the average realized data rate within a system, it characterizes the performance of an actual radio access technology with a certain choice of RRM schemes and system parameters. A different concept is the notion of *channel capacity* which constitutes the maximum rate with which information can be transmitted error-free over a given communication channel.

Classical information-theoretic research aims at establishing bounds on the information rate of single links (i.e., between one transmitter and one receiver) and devises communication schemes to achieve these bounds. It also considers network settings where multiple links are operated simultaneously and interfere with each other. Questions in that context are, for example, how transmission on multiple links can be optimized jointly. Usually, this involves a trade-off between different links so that instead of a single capacity value, a rate-region as the combination of jointly feasible rates is considered.

In the following sections, we will briefly touch on some information-theoretic capacity measures as we introduce basic communication schemes like *multiple input multiple output* (MIMO) communications. Research on fundamental capacity bounds and means to achieve them is important as it can direct the implementation and standardization of

radio access technologies towards higher spectral efficiencies. In contrast to the modeling assumptions under which these results are usually derived, the RRM schemes for actual radio access technologies like LTE have to operate with limited signaling and computational resources. Their performance in the field will therefore often be significantly below theoretical bounds. In Section 2.3, we will discuss some of the reasons for that. The basic and advanced RRM schemes we will discuss in Chapters 5 and 6 aim at improving the spectral efficiency but they are not an *optimization* in a mathematical sense: While an objective function might be easy to define (e.g., the weighted sum rate over all links), the various constraints imposed by actual systems already make an optimization problem definition difficult. For example, actual systems have to operate with incomplete and delayed information, obey their users' *quality of service* (QoS) requirements, and operate within the limits of standards.

2.1.1. System Model for a Single Wireless MIMO Link

The Radio Channel in Complex Baseband Representation

The communication channels considered for cellular mobile networks are, of course, radio channels, where an electromagnetic radio wave propagates from the transmitter's antenna to the receiver's antenna. The transmission leaving the transmitter's antenna is a time-varying, real-valued radio frequency signal $x_{\text{RF}}(t)$ that describes, for example, the amplitude of the electrical field. To operate within the assigned radio spectrum, the information is modulated onto a carrier frequency f_c :

$$x_{\text{RF}}(t) = \Re \left\{ x(t) \cdot e^{2\pi j f_c t} \right\} \quad (2.1)$$

$$= \Re \left\{ [x_I(t) + j \cdot x_Q(t)] \cdot e^{2\pi j f_c t} \right\} \quad (2.2)$$

$$= x_I(t) \cos(2\pi f_c t) - x_Q(t) \sin(2\pi f_c t). \quad (2.3)$$

Here, $x(t) = x_I(t) + j \cdot x_Q(t)$ is called the *complex baseband representation* of the transmitted signal. The $x_I(t)$ part is called the *in-phase* component and the $x_Q(t)$ part is called the *quadrature* component. Both directly modulate two sinusoid carriers at the center frequency f_c that are shifted against each other by 90° , see (2.3). When propagating from the transmitter to the receiver via a certain propagation path denoted with an index i , the radio signal experiences a *propagation delay* τ_i and only arrives with a fraction β_i of its original power. We can thus express the amplitude of the received radio signal on path i as [Ivr05]

$$y_{\text{RF}}(t) = \sqrt{\beta_i} \cdot x_{\text{RF}}(t - \tau_i) \quad (2.4)$$

$$= \Re \left\{ \sqrt{\beta_i} \cdot x(t - \tau_i) \cdot e^{2\pi j f_c (t - \tau_i)} \right\} \quad (2.5)$$

$$= \Re \left\{ \sqrt{\beta_i} \cdot e^{-2\pi j f_c \tau_i} \cdot x(t - \tau_i) \cdot e^{2\pi j f_c t} \right\}. \quad (2.6)$$

In the following, we will stick to the complex baseband notation where, by comparing (2.6) with (2.1), we can use the *complex path coefficient* c_i for the considered path

$$c_i = \sqrt{\beta_i} \cdot e^{-2\pi j f c \tau_i} \quad (2.7)$$

to express channel's influence along the considered propagation path i on the phase and the amplitude of the received signal in complex baseband representation:

$$y(t) = c_i \cdot x(t - \tau_i). \quad (2.8)$$

In the wireless channels encountered in typical cellular communications settings, the transmitter's signal reaches the receiver via multiple different paths as described in the next section. With the help of the Dirac delta function, we can incorporate the different propagation delays τ_i into the so-called *channel impulse response* (CIR) $c(t)$ summarizing all N_{MPC} propagation paths (called *multipath components*):

$$c(t) = \sum_{i=1}^{N_{\text{MPC}}} c_i(t) \cdot \delta(t - \tau_i). \quad (2.9)$$

In the following, we are interested in the frequency-domain representation of the channel, which is obtained by means of a discrete Fourier transform (DFT) from the complex baseband I/Q samples of $x_I(t)$ and $x_Q(t)$. After the DFT we obtain the complex *channel transfer function* (CTF) for some time instant t and some frequency f :

$$h(t, f) = \sum_{i=1}^{N_{\text{MPC}}} c_i(t) \cdot \exp(-2\pi j f \tau_i). \quad (2.10)$$

Single Link MIMO System Model

When we introduced the complex baseband representation above, we implicitly assumed a single antenna at both ends of the link. Modern wireless systems, however, usually employ multiple antennas at both ends. Figure 2.1 shows the common classification of multi-antenna systems, where the number of inputs to the channel as seen from the transmitter (on the right) and the number of outputs from the channel as seen from the receiver (on the left) are used to distinguish *single input single output* (SISO), *single input multiple output* (SIMO), *multiple input single output* (MISO), and *multiple input multiple output* (MIMO) schemes. In order to accommodate all these possibilities, we thus denote the inputs and outputs to the channel by complex vectors $\mathbf{x} \in \mathbb{C}^{N_{\text{Tx}}}$ and $\mathbf{y} \in \mathbb{C}^{N_{\text{Rx}}}$, which are sized according to the number N_{Rx} of receive (Rx) and the number N_{Tx} of transmit (Tx) antennas, respectively. In accordance with the input and output sizes, the channel behavior is described by a scalar $h \in \mathbb{C}$ for the SISO case, vectors $\mathbf{h} \in \mathbb{C}^{N_{\text{Rx}} \times 1}$ for the SIMO and $\mathbf{h}^T \in \mathbb{C}^{1 \times N_{\text{Tx}}}$ for the MISO case, and by a matrix $\mathbf{H} \in \mathbb{C}^{N_{\text{Rx}} \times N_{\text{Tx}}}$ for the MIMO case. Figure 2.1 also shows that the received signal is impaired by an additive white complex baseband noise $\mathbf{n} \in \mathbb{C}^{N_{\text{Rx}}}$. The noise is modeled as a random vector whose entries per Rx antenna are zero mean circularly symmetric

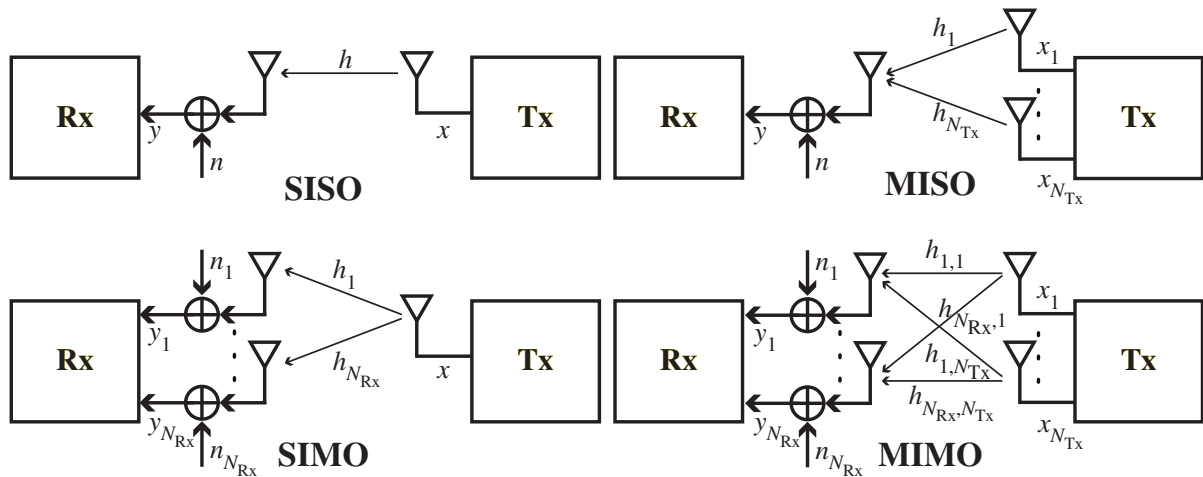


Figure 2.1.: Classification of MIMO schemes

Gaussian random variables [Gal08] denoted by $[n]_i \sim \mathcal{CN}(0, \sigma_{\text{th}}^2)$ where the variance σ_{th}^2 is the noise power per Rx antenna. Here, we assume that the noise results from temperature-dependent thermal noise (-174 dBm/Hz at room temperature) in the receiver components. It is white both in the frequency domain (i.e., constant over the relevant frequency range) and in the spatial domain, meaning that noise samples between antennas are uncorrelated with $E[\mathbf{nn}^H] = \sigma_{\text{th}}^2 \mathbf{I}$. This allows us to describe the transmission as:

$$\begin{aligned} \mathbf{y} &= \mathbf{H}\mathbf{x} + \mathbf{n} \\ &= \begin{pmatrix} h_{1,1} & \cdots & h_{1,N_{\text{Tx}}} \\ \vdots & \ddots & \vdots \\ h_{N_{\text{Rx}},1} & \cdots & h_{N_{\text{Rx}},N_{\text{Tx}}} \end{pmatrix} \begin{pmatrix} x_1 \\ \vdots \\ x_{N_{\text{Tx}}} \end{pmatrix} + \begin{pmatrix} n_1 \\ \vdots \\ n_{N_{\text{Rx}}} \end{pmatrix}. \end{aligned} \quad (2.11)$$

Fading Channel Coefficients and their Correlations

As already mentioned above, the complex coefficients $c_i(t)$ and thus the frequency domain representations $h(t, f)$ capture all propagation effects encountered in the wireless channel between transmitter and receiver. The literature distinguishes three main effects that influence the channel gain $|h(t, f)|^2$. First, the dominant large-scale effect is the so-called *pathloss*. In free space, the power emitted into a certain solid angle by the transmitter distributes onto a surface segment whose area grows with increasing distance d so that the received power is attenuated accordingly proportional to $(1/d)^2$. More general pathloss models account for other large-scale effects like additional attenuation from buildings etc. so that often pathloss exponents > 2 are encountered. The second component affecting the channel gain is called the *shadowing* or *slow fading*, which represents the signal attenuation or gain from larger objects like buildings. It is often modeled as a log-normal component adding to the pathloss. The term *slow fading*

is used because the attenuation slowly varies as a user travels, for example, along a street. This is in contrast to the third classical effect called the *fast fading*. This kind of fading varies much faster because it results from the user's relative velocity with respect to the transmitter or surrounding objects that create a multipath interference field due to reflection, diffraction, and scattering effects. Constructive and destructive interference locations are spaced apart on the order of wavelengths (decimeters for the relevant frequency range), which is a distance that even a slowly moving user passes within short time.

If there is no direct *line-of-sight* (LoS) connection between transmitter and receiver, the individual multipath components have similar average powers and i.i.d. Gaussian-distributed in-phase and quadrature components of their channel coefficients in complex baseband representation. The power of the resulting superposition $h(t, f)$ is then *Rayleigh*-distributed. We thus speak of *Rayleigh fading*, which is characterized by deep fades that occur when the multipaths (almost) cancel each other. If there is an LoS ray, it is usually much stronger than the other multipaths so that a complete extinction of the signal is impossible. This fading situation is best modeled by a *Rician* distribution.

The channel coefficients $h(t, f)$ evolve both in the time and frequency domain. The *coherence time* measures how fast a channel changes over time. There are multiple definitions in the literature [Gal08], depending on the model and on the threshold used to decide if the channel coefficient has changed enough to be considered different. Here, we define the coherence time T_c for which the channel stays roughly constant as [Gal08]:

$$T_c := \frac{1}{2 \cdot \sigma_D}, \quad (2.12)$$

where $\sigma_D = \max_i \Delta f_i - \min_i \Delta f_i$ is the *Doppler spread* denoting the difference between the maximum and minimum *Doppler shift* $\Delta f_i = \frac{\Delta v_i}{c} \cdot f_c$. The Doppler shift results from the user's relative velocities Δv_i with respect to all multipaths where c denotes the speed of light. If we consider a scenario where a user moves with speed v towards a reflecting wall, we have a Doppler spread $\sigma_D = 2 \cdot \Delta f = \frac{2 \cdot v}{c} \cdot f_c$ and thus a coherence time of

$$T_c = \frac{c}{4 \cdot v \cdot f_c}. \quad (2.13)$$

The situation is very similar in the frequency domain. The *coherence bandwidth* indicates over which frequency bandwidth the channel changes only slightly. It is inversely proportional to the *delay spread* σ_τ which we define [Gal08] as the difference $\sigma_\tau = \max_i \tau_i - \min_i \tau_i$ between the propagation delay τ_i on the shortest and longest multipath. Based on that, we use the following definition of the coherence bandwidth B_c from the literature [Gal08]:

$$B_c := \frac{1}{2 \cdot \sigma_\tau}. \quad (2.14)$$

If the coherence bandwidth that results from the multipath propagation is larger than the frequency bandwidth of a wireless channel, we call the channel *frequency-flat*. Otherwise, we say that the channel is *frequency-selective*. For the capacity considerations

in the following section, we will assume that a single channel is so *narrowband* that it is frequency-flat. Also, we will assume that the time per channel use is shorter than the coherence time so that the channel is constant per channel use. Between different channel uses the channel might change, which is called *block-fading*.

The wireless channel with multipath propagation does not only exhibit correlated fading in the time and frequency dimension but also in the spatial dimension. In fact, the coherence time results from the relative movement between the antenna and the multipath interference field. For a transmitter or receiver with multiple antennas, this means that channel samples taken at different antenna positions are not necessarily identical but do exhibit some kind of *spatial correlation*. In the following, we will discuss the situation from the receiver's point of view, but the situation is identical for the transmitter. If the different multipaths in a *non-line-of-sight* (NLoS) setting arrive from a wide range of angles, or are even uniformly distributed around the receiver as in Jake's model [Jak74; Gol05], the relative phases between the multipaths are different at spatially separated antenna locations leading to a different Rayleigh fading behavior. In the case of uniformly distributed angles, the signals at two antennas experience uncorrelated fading at an antenna spacing of about $0.4 \cdot \lambda$ [Gol05]. In the other extreme, we consider the case where we only have one single (LoS) propagation path. We assume that the distance between the Tx and Rx antenna arrays is large compared to the antenna separation in the array and the wavelength [GBGP02] so that we can approximate the arriving waves by parallel wavefronts. Then, we have a fixed relationship between two separate antennas that depends on the phase shift φ_i that the incoming wave front experiences between impinging on the closest antenna (antenna 1) and some other farther antenna $i \neq 1$. Assuming that the additional pathloss between the antennas in the antenna array is negligible, the complex path coefficient (see (2.7)) for antenna i is $c_i = \exp(j\varphi_i) \cdot c_1$. The same holds true for the transmitter antennas where we have $c_j = \exp(j\varphi_j) \cdot c_1$ with respect to the first Tx antenna. That means that in the considered situation with a single LoS ray, we have a frequency-domain MIMO matrix H where each row $i \neq 1$ (Rx side) linearly depends on the first row, and where each column $j \neq 1$ (Tx side) linearly depends on the first column. Thus, such a MIMO channel matrix has *rank* one and we say that the MIMO channel has rank one in this case. If, however, there are many rays with a high *angular spread* and if the antennas are sufficiently spaced apart on both sides of the link, the fading between antenna pairs is expected to have low correlation. In that case the MIMO channel matrix can have full rank. In Chapter 3 we will discuss in detail how the typical MIMO channels in cellular deployment scenarios are modeled by the IMT-Advanced channel model.

2.1.2. Single Link MIMO Channel Capacity

The channel or *Shannon* capacity is the maximum information rate, also called *mutual information*, that can be communicated error-free in the long term over a given noisy channel. Shannon showed [Sha48] that the maximum mutual information is reached if the input x to the channel is Gaussian distributed. At the same time, assuming the

noise to be Gaussian distributed (with a variance σ_{th}^2 as the noise power) is the worst-case assumption for the noise distribution [CT06]. For the additive white Gaussian noise (AWGN) SISO channel introduced in (2.11) the capacity is given by [Sha48]

$$C = B \log_2 \left(1 + \frac{|h|^2 P_{\text{tx}}}{\sigma_{\text{th}}^2} \right) \quad \text{bit/s}, \quad (2.15)$$


where σ_{th}^2 is the noise power in the considered bandwidth B , and $P_{\text{tx}} = E[|x|^2]$ is the average transmit power of the complex-baseband input symbols x .

In this most basic and famous Shannon capacity expression, the channel is considered *time-invariant*. That is, the channel gain h is fixed and only the instantaneous noise varies according to a normal distribution with variance σ_{th}^2 . In actual wireless systems, the channel is usually both *frequency-selective* and *time-variant*. This gives rise to various capacity definitions, which also depend on the amount of knowledge about the current channel state. For example, if the transmitter knows the fading status of different subcarriers and time slots, it can distribute its power in a so-called *water-filling* way to the most favorable resources resulting in a higher overall capacity compared to an average capacity with equal power loading. Such an average capacity is called an *ergodic capacity*, which is defined as the average of all faded capacities weighted with their probability of occurrence. Another traditional measure is the *outage capacity* [Gol05], which gives the maximum feasible rate at which at most a certain percentage of fading channel realizations does not admit transmission with a sufficient rate. In the following, we will concentrate on the *ergodic capacity*. We will not consider *outage capacity* because with modern RRM methods in cellular systems, ergodic capacity is the more relevant measure [JL10].

In modern cellular systems, the SISO channel is not very important because base stations and terminals usually are equipped with multiple antennas so that we have a MIMO configuration. For the MIMO channel, Telatar [Tel99] first gave an expression for the mutual information between the output $\mathbf{y} \in \mathbb{C}^{N_{\text{Rx}}}$ at the receiver and the channel input $\mathbf{x} \in \mathbb{C}^{N_{\text{Tx}}}$ at the transmitter. For the general case of a spatially-correlated noise vector \mathbf{n} at the receiver the mutual information MI is given by (see, e.g., [Ivr05]):

$$\text{MI}(\mathbf{x}, \mathbf{y}) = B \log_2 \det \left(\mathbf{I}_{N_{\text{Rx}}} + \mathbf{H} \mathbf{Q} \mathbf{H}^H \mathbf{R}_n^{-1} \right) \quad \text{bit/s}, \quad (2.16)$$

where $\mathbf{R}_n = E\{\mathbf{n}\mathbf{n}^H\}$ and $\mathbf{Q} = E\{\mathbf{x}\mathbf{x}^H\}$ are the covariance matrices of the received noise and the transmitted signal, respectively. Given the channel \mathbf{H} and noise conditions \mathbf{R}_n , we can obtain the capacity of the MIMO channel by maximizing the mutual information in (2.16) with respect to the only remaining variable \mathbf{Q} . As we will see, the degree to which we can maximize the mutual information depends on how much *channel state information* (CSI) the transmitter has in order to choose a suitable transmit covariance \mathbf{Q} . In the downlink, the receiver can easily obtain receiver-side channel state information (CSIR) from pilot signals. These are routinely transmitted in modern systems so that we assume that the receiver knows the channel, that is, it has perfect CSIR. Providing channel knowledge to the transmitter (CSIT) is more


 Figure 2.2.: MIMO system model for a single link with precoder \mathbf{P} and Rx filter \mathbf{W}^H

difficult unless transmitter and receiver operate in a time-division duplexing (TDD) fashion on the same frequency bandwidth so that the transmitter can exploit channel reciprocity. Otherwise, the receiver has to provide feedback. To influence the transmit covariance $\mathbf{Q} = E\{\mathbf{x}\mathbf{x}^H\}$, we assume (without loss of generality [Ivr05]) that the transmitter does not directly modify the complex baseband symbol vector \mathbf{x} transmitted on the antennas. As shown in Fig. 2.2, it instead uses a precoder matrix \mathbf{P} to map a vector $\mathbf{s} \in \mathbb{C}^M$ that contains the actual information symbols to the transmit vector $\mathbf{x} = \sqrt{P_{\text{tx}}}\mathbf{P}\mathbf{s}$ where P_{tx} is the total transmit power. We restrict the precoder to have unit norm $\|\mathbf{P}\|_{\text{F}}^2 = \text{tr}(\mathbf{P}\mathbf{P}^H) = 1$ so that it does not scale the total transmit power. In the following, we will call the product $\mathbf{H}\mathbf{P}$ the *effective channel*. The entries in \mathbf{s} are assumed to be i.i.d. complex normal (i.e., unity power) with $E\{\mathbf{s}\mathbf{s}^H\} = \mathbf{I}_M$ where, for the time being, we will assume that the number of layers M in \mathbf{s} equals the number of Tx antennas N_{Tx} . That means that the transmit covariance $\mathbf{Q} = E\{\mathbf{x}\mathbf{x}^H\} = P_{\text{tx}}\mathbf{P}\mathbf{P}^H$ depends only on the chosen power and the precoder matrix. While treating a single link, we will restrict the discussion to spatially white noise with $\mathbf{R}_n = \sigma_{\text{th}}^2 \mathbf{I}_{N_{\text{Rx}}}$ where σ_{th}^2 is the noise power per Rx antenna. Thus, we can write (2.16) as:

$$\text{MI}(\mathbf{s}, \mathbf{y}) = B \log_2 \det \left(\mathbf{I}_{N_{\text{Rx}}} + \frac{P_{\text{tx}}}{\sigma_{\text{th}}^2} \mathbf{H}\mathbf{P}\mathbf{P}^H \mathbf{H}^H \right) \text{ bit/s.} \quad (2.17)$$

If no CSIT is available, it is optimal [Tel99] to use a scaled identity matrix as the precoder $\mathbf{P} = \sqrt{1/N_{\text{Tx}}}\mathbf{I}_{N_{\text{Tx}}}$ for which $\text{tr}(\mathbf{P}\mathbf{P}^H) = 1$ holds. Thus, the capacity without CSIT (but with perfect CSIR) is:

$$C_{\text{No CSIT}} = B \log_2 \det \left(\mathbf{I}_{N_{\text{Rx}}} + \frac{P_{\text{tx}}}{\sigma_{\text{th}}^2} \mathbf{H}\mathbf{H}^H \right) \text{ bit/s.} \quad (2.18)$$

In the ideal case, the transmitter has perfect and instantaneous knowledge of the channel \mathbf{H} . Using a *singular value decomposition* (SVD) [Str09], we can write any channel matrix as $\mathbf{H} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^H$ where \mathbf{U} and \mathbf{V} are unitary matrices for which $\mathbf{U}\mathbf{U}^H = \mathbf{U}^H\mathbf{U} = \mathbf{I}_{N_{\text{Rx}}}$ and $\mathbf{V}\mathbf{V}^H = \mathbf{V}^H\mathbf{V} = \mathbf{I}_{N_{\text{Tx}}}$ holds. Further, $\mathbf{\Sigma} \in \mathbb{C}^{N_{\text{Rx}} \times N_{\text{Tx}}}$ is a diagonal matrix that contains $S \leq \min(N_{\text{Tx}}, N_{\text{Rx}})$ non-zero so-called *singular values* ζ_1, \dots, ζ_S of \mathbf{H} . Note that the singular values ζ_i of \mathbf{H} are the square roots of the eigenvalues λ_i (i.e., $\zeta_i = \sqrt{\lambda_i}$) of

both $\mathbf{H}\mathbf{H}^H$ and $\mathbf{H}^H\mathbf{H}$ [Str09]. If we choose \mathbf{V} as our precoder ($\mathbf{P} = \mathbf{V}$) and further filter the precoded channel with the matrix \mathbf{U}^H at the receiver (see Fig. 2.2), the precoded and filtered channel

$$\mathbf{U}^H\mathbf{H}\mathbf{V} = \mathbf{U}^H\mathbf{U}\mathbf{\Sigma}\mathbf{V}^H\mathbf{V} = \mathbf{\Sigma} \quad (2.19)$$

decomposes into $S = \text{rank}(\mathbf{H}) \leq \min(N_{\text{Tx}}, N_{\text{Rx}})$ parallel channels that do not interfere with each other. The noise $\mathbf{n}' = \mathbf{U}^H\mathbf{n}$ gets filtered by \mathbf{U}^H , which is unitary and independent of \mathbf{n} and hence does not affect the channel capacity. We can treat the resulting S independent channels as equivalent SISO channels

$$\hat{s}_i = \sqrt{P_i} \cdot \zeta_i \cdot s_i + n', \quad (2.20)$$

where P_i is the fraction of the total power P_{Tx} that is assigned to subchannel i . Applying the Shannon capacity for the SISO case (2.15), we can write the capacity of the MIMO channel with perfect CSIT (and perfect CSIR) as [Gol05]:

$$C_{\text{CSIT}} = B \max_{P_1, \dots, P_S} \sum_{i=1}^S \log_2 \left(1 + \frac{P_i}{\sigma_{\text{th}}^2} \lambda_i \right) \text{ bit/s} \quad \text{with} \quad \sum_{i=1}^S P_i = P_{\text{Tx}}. \quad (2.21)$$

Again, the capacity is achieved by a water-filling power allocation strategy that allocates most power to those *eigenmodes* (columns in the precoder \mathbf{V}) that have the largest singular value and thus equivalent channel gain. For channels with rank one or in low *signal to noise ratio* SNR situations, the optimal waterfilling solution allocates all power to the strongest eigenmode so that only a single signal layer is transmitted using the corresponding column in \mathbf{V} as the beamforming vector (*dominant eigenmode transmission*). Conversely, in very high SNR conditions (i.e., for $P_i/\sigma_{\text{th}}^2 \rightarrow \infty$) with full rank channels, all eigenmodes of the channel are used. The gain of channel-aware precoding and power distribution over the eigenmodes vanishes compared to identity precoding without CSIT because the relative differences between the powers per eigenmode become small.

From the general solution for the MIMO case with perfect CSIT, we can also easily derive the capacities of the SIMO and MISO channels. To simplify the notation, we will denote both the SIMO and the MISO channel with \mathbf{h} here even though the latter would be transposed. For a vector channel, we only get one singular value ζ . It is the square root of the eigenvalue of $\mathbf{h}\mathbf{h}^H$ and $\mathbf{h}^H\mathbf{h}$, which is the scalar $\mathbf{h}^H\mathbf{h} = \|\mathbf{h}\|^2 = \lambda = \zeta^2$. Of course, the whole power P_{Tx} gets allocated to the single channel. Thus, in this case (2.21) becomes

$$C_{\text{CSIT}}^{(\text{MISO})} = C_{\text{CSIR}}^{(\text{SIMO})} = B \log_2 \left(1 + \frac{P_{\text{Tx}}\|\mathbf{h}\|^2}{\sigma_{\text{th}}^2} \right) \text{ bit/s}. \quad (2.22)$$

In the presence of white noise, we can reach the capacity in the SIMO case by using the receive filter $\mathbf{W}^H = \mathbf{h}^H/\|\mathbf{h}\|$ to filter the received vector $\mathbf{y} = \mathbf{h}s + \mathbf{n}$ which gives us a scalar channel with an effective power gain of $\|\mathbf{h}\|^2$ and an unchanged expected noise power:

$$\hat{s} = \mathbf{W}^H\mathbf{y} = \frac{\mathbf{h}^H}{\|\mathbf{h}\|} (\mathbf{h}s + \mathbf{n}) = \|\mathbf{h}\|s + \frac{\mathbf{h}^H}{\|\mathbf{h}\|}\mathbf{n}. \quad (2.23)$$

We will call this scheme maximum ratio combining (MRC) in later chapters. In the MISO case, we can reach the capacity with perfect knowledge of $\mathbf{h} \in \mathbb{C}^{1 \times N_{\text{Tx}}}$ by beamforming with the precoder vector $\mathbf{p} = \mathbf{h}^H / \|\mathbf{h}\|$ into the direction of the channel. The user effectively sees a scalar channel $y = \mathbf{h}\mathbf{h}^H / \|\mathbf{h}\| \cdot s + n$ with a power gain of

$$\frac{\mathbf{h}\mathbf{h}^H (\mathbf{h}\mathbf{h}^H)^*}{\|\mathbf{h}\| \|\mathbf{h}\|} = \frac{\|\mathbf{h}\|^2 \cdot \|\mathbf{h}\|^2}{\|\mathbf{h}\|^2} = \|\mathbf{h}\|^2. \quad (2.24)$$

We have already discussed two different MIMO schemes: in (2.21) we have shown that with multiple antennas at both sides of the link, we can obtain a so-called *spatial multiplexing* gain of $S = \text{rank}(\mathbf{H})$, where the capacity can increase by up to a factor of $S \leq \min(N_{\text{Tx}}, N_{\text{Rx}})$ compared to the SISO case. In general, the same scaling factor also applies to the case without channel state information at the transmitter (CSIT) in (2.18) even though with CSIT we can achieve a higher capacity with the waterfilling power allocation. If the channel has rank one, which is naturally also the case for both SIMO and MISO configurations, we cannot achieve a spatial multiplexing gain but we can achieve an *array* or *beamforming* gain as shown in equations (2.22) and (2.24). A third class of MIMO transmission schemes aims at a *diversity* gain that can be obtained if the fading between antenna pairs is uncorrelated. The more independently fading paths between antenna pairs exist, the lower is the probability that all of them are in a bad fading situation simultaneously. In Chapter 4 we will introduce the most famous diversity scheme proposed by Alamouti [Ala98].

The antenna configuration at both sides of the link, the correlation and fading properties of the propagation channel \mathbf{H} , and the available channel knowledge determine what MIMO scheme is most beneficial for the indented use case. This motivates the need for an accurate MIMO channel model as the one we will introduce in Chapter 3.

2.1.3. From Single Links to Multi-Cell Systems

Multiple Users per Cell: Orthogonal Multiple Access Schemes

So far, we have only considered a single link between one transmitter and one receiver. In an actual cellular system there are, of course, multiple users that want to transfer data simultaneously to and from their serving base station. They thus have to share the spectral resources. Cellular systems have a centralized organization where the *scheduling* mechanism in the base station is in charge of controlling the access to the channel. The two most basic multiple access schemes are *time-division multiple access* (TDMA) and *frequency-division multiple access* (FDMA), where different users are assigned orthogonal resources in the time and frequency domain, respectively. TDMA and FDMA have been used in second generation cellular systems like GSM and they are also present in fourth generation LTE systems in the form of *orthogonal frequency-division multiple access* (OFDMA) as shown later in this chapter. Third generation systems, for example UMTS, use *code-division multiple access* (CDMA), which separates users by applying orthogonal codes. We will not further discuss CDMA in this thesis.

Serving users on orthogonal resources with the help of TDMA and FDMA yields separate channels whose capacities can be modeled as shown in the previous section. The only difference is that the individual channels have less bandwidth (FDMA) or are only available for a fraction of the time (TDMA). Being able to choose from a set of users with independent channels when assigning time and frequency resources gives the scheduler an additional degree of freedom called *multi-user diversity*. If the scheduler has channel state information with respect to the time- and frequency-selective fading of the individual channels, it can schedule transmissions to or from users on those parts of the channel where the user perceives the best channel conditions. That way, a significant gain can be obtained, which we will evaluate in the context of LTE in Chapter 5.

Non-Orthogonal Channel Access with Intra-Cell Interference

A different scheduling strategy is to serve multiple users on the same time and frequency resources. If no further measures are taken, the concurrent transmissions *interfere* at the base station (uplink) or at the user terminals (downlink). That is, the receiver receives a superposition of its desired signal and of all other signals belonging to concurrent transmissions on the same frequency bandwidth. First, we discuss *intra-cell interference*, which is caused by interfering transmissions that happen in the same cell, that is, to and from the same base station. Later in this section, we will discuss *inter-cell interference* and extend the system model from (2.11) to take intra- and inter-cell interference into account, see (2.25).

There are two fundamental ways how a receiver can deal with interference impacting its desired transmission. On the one hand, it can treat the interference as another source of (complex) Gaussian noise in addition to the thermal noise that we have considered so far. The assumption of a Gaussian distribution is justified by the central limit theorem if arbitrary interfering signals from many sources add up. But also interference from a single source can be considered as Gaussian noise because, as discussed in Section 2.1.2, any transmitter would want to use (close to) Gaussian-distributed input signals because that is optimal for the achievable rate. If interference is treated as noise, its power adds to the noise power and instead of the SNR we use the *signal to interference and noise ratio* (SINR) to characterize the achievable rates. Thus, interference is always detrimental to the capacity in this model.

On the other hand, if the interference that adds to the desired signal is known at the receiver, it can be subtracted (canceled) from the sum signal so that the desired signal can be received interference-free. Of course, the exact interfering signal is usually not known at the receiver in advance. But if the receiver has enough knowledge about the structure (i.e., the effective channel, modulation and coding scheme as introduced below) of the interference, it can employ a technique called *successive interference cancellation* (SIC). SIC successively detects all concurrent transmissions, starting with the strongest. Once a transmission has been decoded, and after estimating the interfering channel, the interfering signal can be reconstructed and subtracted (canceled) from the remaining signal to allow the detection of the next strongest transmission without the

interference from previously canceled transmissions. Due to the requirement that the receiver needs to be aware of the interfering transmissions' structure in order to apply SIC, it is most often used in the uplink, where a base station knows in advance about arriving transmissions because it had previously scheduled them. As we focus on the downlink in this thesis, we will not further consider any interference cancellation schemes in the following.

A converse strategy called *dirty paper coding* (DPC) going back to a seminal paper by Costa [Cos83] can be applied at the transmitter. It exploits the fact that noise (or interference in this case) that superimposes with the desired signal at the receiver is detrimental to the capacity only if it is not known to the transmitter before the transmission. In case the transmitter perfectly knows all channels and interfering signals in advance, it can encode its transmission in a way that it offsets the noise (the "dirt") at the receiver. If the interfering transmissions all stem from one transmitter or from transmitters that jointly coordinate their transmissions, it is possible to sequentially apply DPC when encoding the transmissions. That way, a transmission can be precoded to take previously encoded transmissions into account so that their interference vanishes at the intended receiver and only the interference of subsequently encoded transmissions remains. Dirty paper coding is an important concept in information theory because it has been shown [WSS06] that the DPC rate region coincides with the Gaussian MIMO broadcast channel capacity region. The notion of a *broadcast channel* [CT06] applies to the downlink of a cell in which a base station transmits separate data streams to multiple users using the same time and frequency resources. Compared to a channel capacity that indicates the maximum feasible rate on a single link, a capacity region characterizes the combination of rates to different users that are feasible at the same time. Even though DPC is capacity-achieving, it is merely a theoretical concept because in addition to a high computational complexity, it needs non-causal information (i.e. information about future conditions) about interference (from other sources) at the transmitter as well as perfect channel knowledge of all channels [HPV12]. For that reason, DPC plays no role in actual systems and we will not further consider it here.

Finally, another way to (quasi) orthogonalize transmissions to or from different users that share the same time/frequency resource is to apply MIMO processing at one or both ends of the link. A multi-antenna receiver can accept interference as Gaussian noise but can exploit that, in contrast to thermal noise, interference usually does not appear spatially white at the receiver. This allows to combine the Rx signals in a way that power from interfering channel directions is suppressed. A similar concept can be applied at the transmitter with CSIT. For example, by means of a so-called *zero-forcing* precoding scheme, the transmitter can precode two transmissions to two distinct users in a way that no interference power reaches the unintended user. Such schemes are sometimes called *space-division multiple access* or more general, and more popular recently, *multi-user MIMO* (MU-MIMO) transmission. We will discuss MU-MIMO schemes in Chapter 6.

Multi-Cell Systems

While sharing the spectral resources between multiple users offers multi-user diversity, it also means that a user gets a smaller resource share and thus lower rates if there are more and more users who need to share. Also, with a single base station users can be far away from the base station so that they have low-capacity channels due to a high pathloss. The solution to this problem comes in the form of a multi-cell deployment. The basic cellular principle, which is still in use today, was already conceived in the 1940s at the Bell Labs [Rin47]. A geographical area is partitioned into a number of cells, where the users in one cell are served by a dedicated base station as already shown in Fig. 1.2. Having many small cells, as opposed to just one big coverage area, reduces the average distance between BS and MS and thus allows better signal to noise ratios. It also allows to reuse identical frequencies in different cells in the system. In that case, transmissions are not only impacted by thermal noise and intra-cell interference, but also by *inter-cell interference* resulting from transmissions on the same time/frequency resources (*co-channel*) in neighboring cells. To reduce the interference impact of co-channel transmissions, it was suggested early on [Rin47; Mac79] to divide the system frequency bandwidth in subbands so that adjacent cells can operate on different bands. The advantage of such a *frequency reuse* planning is that the same frequency band is only used in a certain *reuse distance*, which ensures that the interfering transmitter is farther away than the serving transmitter. That way, the interfering signal is weaker than the serving signal due to the additional pathloss. Frequency planning has been a very active area of research and different optimization approaches have been applied, for example, to optimize the channel allocation in second generation GSM systems [AvHK+07]. While a limited frequency reuse reduces interference and thus improves SINRs and eventually feasible rates on an inter-cell interference-protected frequency band, it also reduces the availability of spectrum per cell. For that reason, most LTE deployments use a full frequency reuse (also called *reuse 1*) that makes the whole system bandwidth available in all cells. We will discuss this fundamental interference–spectrum trade-off in more detail in Chapter 6.

Figure 2.3 on page 23 shows a schematic view of the intra- and inter-cell interference on some common time/frequency resource in a multi-cell systems for the downlink and uplink direction. We will explain the indexing scheme in the next section when we detail the multi-cell system model. For now, we focus on the general situation. In Fig. 2.3 there are multiple mobiles per cell and we indicate all desired serving signals by black arrows. In order to simplify the illustration, we only show intra- and inter-cell interference (gray arrows) that affects the users in cell number four. The black-colored mobiles are scheduled on the considered resource, whereas the gray-colored mobiles are served at a different time or on a different frequency channel. In the downlink, users three and four are served in a MU-MIMO fashion with different precoders \mathbf{P}_{k_3} and \mathbf{P}_{k_4} so that they receive a desired as well as an interfering signal from their serving base station. In addition, the base stations in cell one and three serve one user each and both of these transmissions cause inter-cell interference to the considered users in cell four. In the uplink, the situation is similar. But now the transmissions

originate from the mobiles, which we assume to only use a single antenna and hence no spatial precoding for the transmission to the serving base station. Base station four receives two simultaneous transmissions from the users in its cell that cause intra-cell interference to each other. In addition, the users in cells one and three cause inter-cell interference to BS four.

Note that there is in general a significant difference in the dynamics of the inter-cell interference situation in the downlink and uplink directions due to the dynamic user-selection by the base station scheduler. In the downlink, the inter-cell interference affecting the users in cell four does not depend on which of their users the neighboring cells schedule (assuming that the precoders in cells one and three do not beamform into a certain direction). In the uplink, in contrast, the inter-cell interference experienced at the BS serving cell four significantly depends on the user-selection in the neighboring cell. If instead of users one and six, which are far away from BS four, users two and five (who are located near the cell border to cell four) transmit, BS four would see a significantly higher inter-cell interference because of the much lower pathloss.

A characteristic feature of cellular systems is the need for a *handover* mechanism that allows to dynamically switch the association of a user from one BS to another as the user moves across cell borders. Such mechanisms would also fall into the category of RRM schemes, but we do not discuss this mobility aspect in this thesis and rather focus on users that stay within their serving cell.

System Model and Mutual Information with Interference for Multi-User and Multi-Cell Systems

We now extend the simple system model from Section 2.1.1 to account for intra- and inter-cell interference. We take a downlink-centric perspective for the system model as this is the primary focus of the thesis. We denote the set of all users (i.e., receivers) in the system as \mathcal{K} with cardinality $N_{\text{MS}} = |\mathcal{K}|$. The set of base stations (i.e., transmitters) is represented by \mathcal{T} with cardinality $N_{\text{BS}} = |\mathcal{T}|$. A certain user k' has exactly one serving BS t' , for which we define a lookup function $t(k') = t'$. The channel from a BS t to a user k is denoted by \mathbf{H}_{tk} . With MU-MIMO it is possible that a BS transmits to more than one user on a given time/frequency resource. We denote the set of users co-scheduled by BS t as \mathcal{D}_t with cardinality $N_{\text{MU},t} = |\mathcal{D}_t|$. With these notational conventions, we can extend the system model from (2.11) to the multi-user and multi-cell case taking arbitrary linear precoders at the base stations into account:

$$\mathbf{y}_{k'} = \underbrace{\sqrt{P_{\text{tx},k'}} \mathbf{H}_{t'k'} \mathbf{P}_{k'} \mathbf{s}_{k'}}_{\text{desired signal}} + \underbrace{\sum_{k \in \mathcal{D}_{t(k')}} \sqrt{P_{\text{tx},k}} \mathbf{H}_{t'k'} \mathbf{P}_k \mathbf{s}_k}_{\text{intra-cell interference}} + \underbrace{\sum_{t \in \mathcal{T} \setminus \{t(k')\}} \sum_{k \in \mathcal{D}_t} \sqrt{P_{\text{tx},k}} \mathbf{H}_{tk'} \mathbf{P}_k \mathbf{s}_k}_{\text{inter-cell interference}} + \mathbf{n}_{k'}. \quad (2.25)$$

On some time/frequency resource, a user k' receives its multi-layer desired transmission $\mathbf{s}_{k'}$ that is precoded with $\mathbf{P}_{k'}$ and sent with Tx power $P_{\text{tx},k'}$ by its serving base

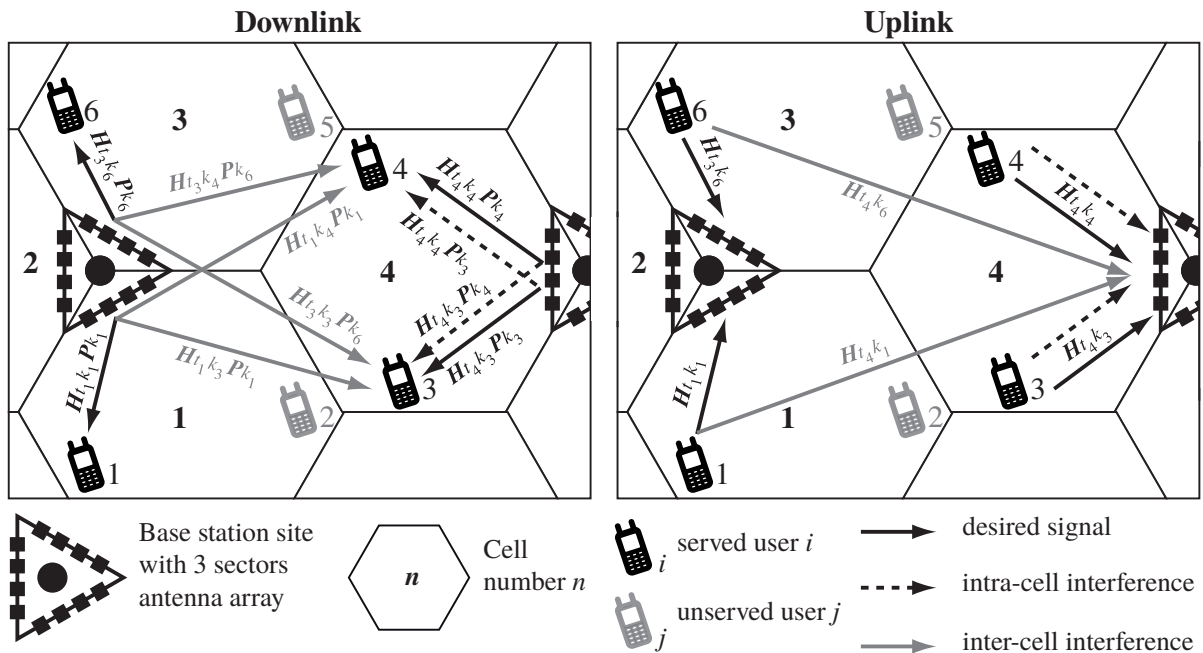


Figure 2.3.: Downlink and uplink intra-cell and inter-cell interference situation (full interference only shown for cell 4)

station $t' = t(k')$, see the examples in Fig. 2.3. It also receives intra-cell interference from all other users $k \in \mathcal{D}_{t'} \setminus \{k'\}$ co-scheduled by BS t' in the same cell as well as inter-cell interference from all users co-scheduled in other cells in addition to noise $\mathbf{n}_{k'}$, which is still assumed to be spatially white. The total *interference plus noise* (IPN) covariance \mathbf{R}_{IPN} is now colored according to the covariances of the interfering signals where we again assume all symbols in all \mathbf{s}_k to be i.i.d. complex normal

$$\mathbf{R}_{\text{IPN},k'} = \sum_{k \in \mathcal{D}_{t(k')} \setminus \{k'\}} P_{\text{tx},k} \mathbf{H}_{t'k'} \mathbf{P}_k \mathbf{P}_k^H \mathbf{H}_{t'k'}^H + \sum_{t \in \mathcal{T} \setminus \{t(k')\}} \sum_{k \in \mathcal{D}_t} P_{\text{tx},k} \mathbf{H}_{tk'} \mathbf{P}_k \mathbf{P}_k^H \mathbf{H}_{tk'}^H + \sigma_{\text{th},k'}^2 \mathbf{I}_{N_{\text{Rx},k'}}. \quad (2.26)$$

As the MIMO mutual information expression (2.16) was derived [Ivr05] for the general case with colored noise covariance, we can use it as the upper bound on the achievable rate on the link between BS t' and user k' with intra-cell and inter-cell interference. Given the chosen transmit covariance $\mathbf{Q} = P_{\text{tx},k'} \mathbf{P}_{k'} \mathbf{P}_{k'}^H$ and all other transmit covariances, which affect the total interference plus noise covariance $\mathbf{R}_n = \mathbf{R}_{\text{IPN}}$, we can write (2.16) as:

$$\text{MI}(\mathbf{s}_{k'}, \mathbf{y}_{k'}) = B \log_2 \det \left(\mathbf{I}_{N_{\text{Rx}}} + P_{\text{tx},k'} \mathbf{H}_{t'k'} \mathbf{P}_{k'} \mathbf{P}_{k'}^H \mathbf{H}_{t'k'}^H \mathbf{R}_{\text{IPN},k'}^{-1} \right) \text{ bit/s}. \quad (2.27)$$

2.2. LTE System Design

So far, we introduced the terminology and some fundamental rate and capacity considerations for MIMO cellular communications. We now turn our attention to the design

of LTE that incorporates many of the aforementioned concepts. We will begin by giving an overview of the LTE protocol stack. As they are most relevant for our spectral efficiency investigation, we will put special emphasis on the physical layer and the radio resource management mechanisms in LTE.

The treatment in this section cannot be a comprehensive introduction to the LTE and LTE-Advanced standard. There are several books that give a more detailed overview [DPS11; HT11; Cox12; STB11]. The standard itself consists of a range of documents published by the 3GPP. The air interface and radio access network aspects can be found in the documents belonging to the series 36 of the 3GPP specifications [TS36].

2.2.1. Overview of the LTE Protocol Stack

With LTE as the successor to UMTS, both the radio access network component, now called *Evolved Universal Terrestrial Radio Access Network* (E-UTRAN), as well as the core network part, now called *Evolved Packet Core* (EPC), have been re-designed as an all-IP *packet-switched* architecture. Here, we will focus on the radio access network and more specifically on the *user plane* that carries the payload data over the radio channel. In the LTE standard, a base station is called *evolved NodeB* (eNB) and a mobile station is called *user equipment* (UE). We will stick to the terms base station (BS) and mobile station (MS) throughout this thesis.

Figure 2.4 shows the protocol stack encompassing the layer 1 and layer 2 of an LTE base station and a mobile station. In the *downlink* (DL), the BS receives IP data flows for different users, which are forwarded through a tunnel from a router in the core network, called the *packet data network gateway* (P-GW). The BS offers so-called *radio bearers* to distinguish between the data flows of different users but also between data flows with different *quality of service* (QoS) requirements belonging to the same user. On the way to the MS, the data travels first through the layer 2, which consists of the *packet data convergence protocol* (PDCP), *radio link control* (RLC), and *medium access control* (MAC) sublayers. The layer 1 is also called the PHY layer and performs different functions as shown in Fig. 2.4. The radio channel between the BS and MS forms the logical *Uu* radio interface in LTE. In the MS, the incoming data traverses the corresponding peer entities in the protocol stack and is finally delivered to higher layers. Besides this user plane data flow, the *control plane* data flow of the *radio resource control* (RRC) protocol between the BS and the MS takes the same path through the protocol stack.

The PDCP sublayer [36.323] takes care of *robust header compression* (ROHC) as well as ciphering to protect the data on the air interface. The header compression provides different profiles to compress the headers of IP packets as well as the carried transport and application layer protocols (e.g., IP, UDP and RTP). Header compression is especially important for short data packets. For example, with the *voice over LTE* (VoLTE) voice call support for LTE, the uncompressed headers would account for more than 50% of the data rate to be transmitted over the radio interface if no compression was in place [PHH+12].

The RLC sublayer prepares one (without spatial multiplexing) or two (with spatial multiplexing, regardless of how many spatial layers are transmitted in parallel) *transport blocks* of the desired size for transmission on the PHY allocated resources. In the process, it concatenates or segments IP packets belonging to the different radio bearers of one user. The receiving peer entity in the MS reassembles the individual IP packets upon reception of the corresponding data segments. The process is thus also called *segmentation and reassembly* (SAR). Depending on the desired reliability of the transmission, the RLC provides three different modes with different degrees of reliability. The *acknowledged mode* (AM) provides SAR functionality, detects and retransmits missing segments by means of an *automatic repeat request* (ARQ) protocol, and assures in-sequence delivery of IP packets. Both the retransmission and in-sequence delivery mechanisms are important to enable high throughput TCP/IP transmissions so that the AM is typically used for such data radio bearers. The *unacknowledged mode* (UM),

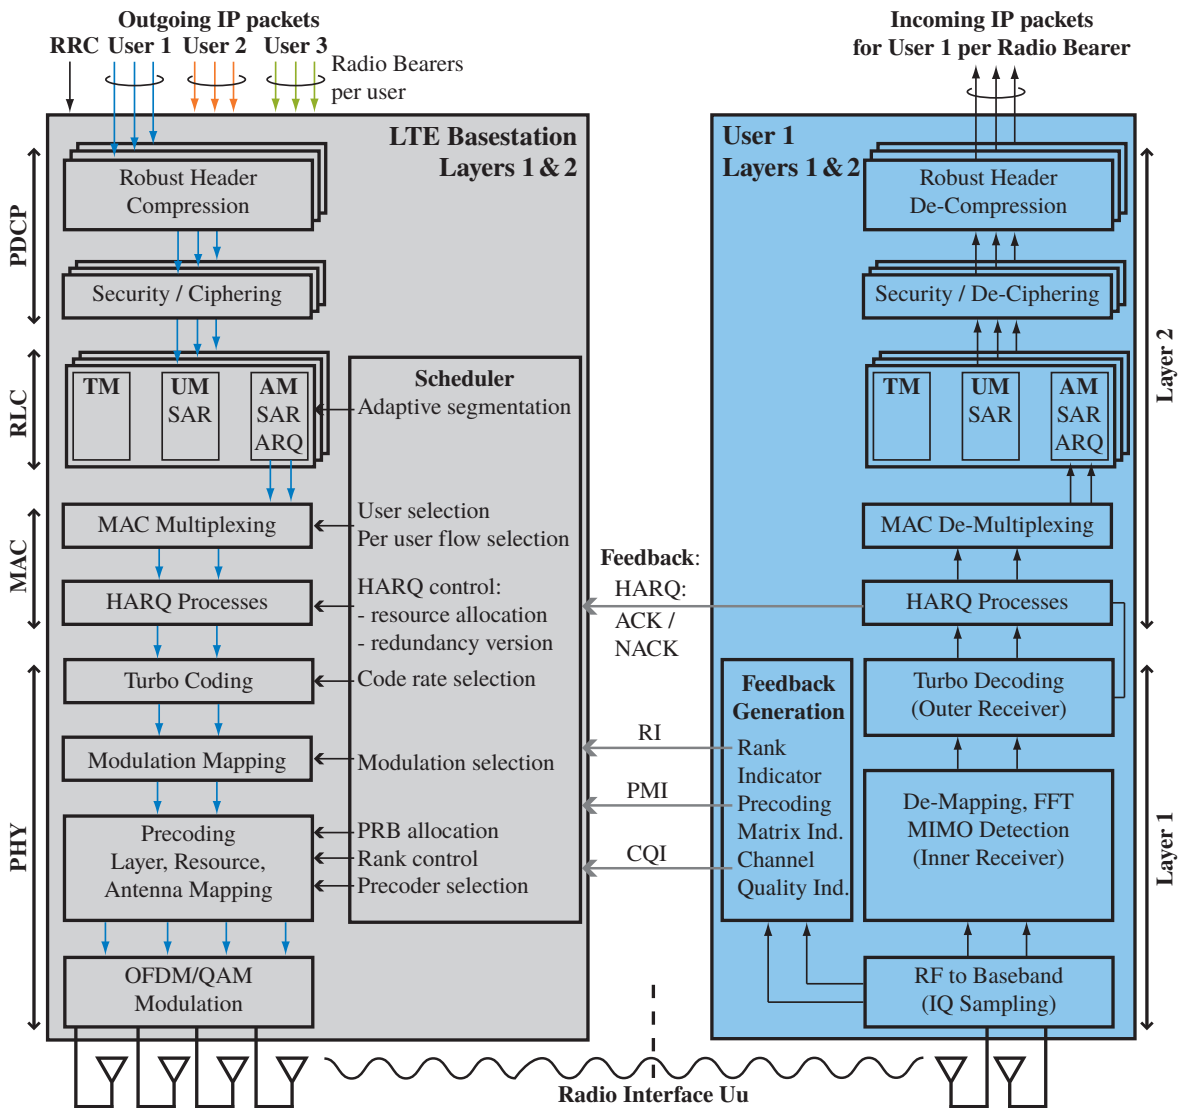


Figure 2.4.: LTE protocol stack layers 1 and 2 with downlink data flow for one user

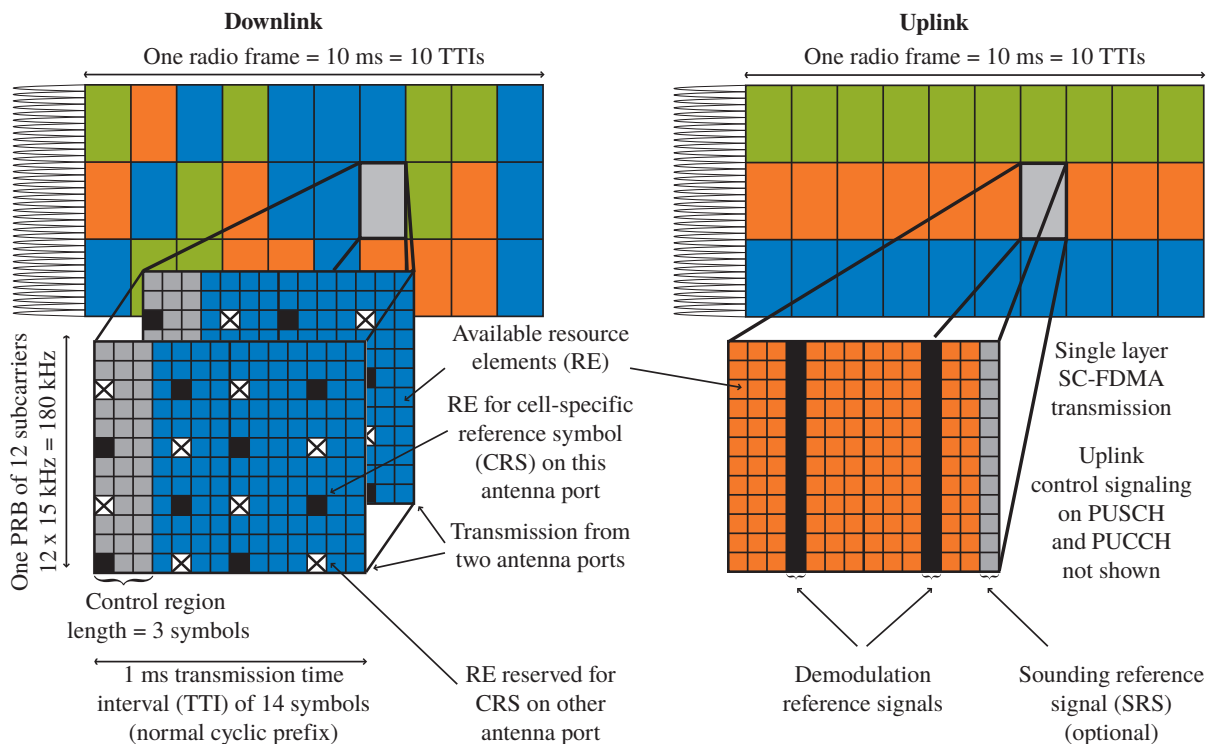


Figure 2.5.: Structure of time and frequency resources for scheduling [36.211]

in contrast, provides SAR functionality and assures in-sequence delivery, but does not retransmit missing segments. Thus, it is used, for example, for VoLTE voice calls where real-time requirements do not allow RLC retransmissions. Finally, the *transparent mode* (TM) transmits packets without any further treatment. It is used for control-plane broadcasts. The MAC sublayer links the layer 2 to the layer 1. According to the standard [36.321], the MAC sublayer is at the heart of the protocol stack because it allocates radio resources, chooses transport formats, and performs data multiplexing as well as *hybrid automatic repeat request* (HARQ) retransmissions. De facto, the scheduling in the LTE downlink is a cross-layer task that spans multiple sublayers from the RLC down to the PHY, as shown in Fig. 2.4.

The base station controls the scheduling for the downlink as well as for the uplink. In the downlink, it can configure the mobile station to report *channel state information* (CSI) feedback in order to adapt the scheduling to the current channel conditions. For the uplink, the BS sends its scheduling decision to the MS, which then performs the transmission accordingly. The uplink data basically takes the same path through the protocol stack as in the downlink but in the opposite direction. We will discuss the scheduling and feedback mechanisms in more detail when we cover the resource allocation schemes in Section 2.2.3. Before that, we will first give a high-level overview of the physical resources that are available to carry the DL and *uplink* (UL) physical channels.

Figure 2.5 shows that the radio resources available to the BS are organized in a grid structure. On the horizontal axis, the time duration of one LTE *radio frame* is shown,

which consists of 10 *transmission time intervals* (TTIs). The scheduling for DL and UL is performed for each 1 ms TTI (also called *subframe* in LTE) by the base station. For each user that is allocated in the downlink, the BS creates one or two transport blocks and assigns frequency resources for the transmission. The smallest allocation in the frequency domain is called a *physical resource block* (PRB) spanning a frequency bandwidth of 180 kHz. In Fig. 2.5 only three PRBs are shown. Depending on the total system bandwidth, between 6 and 100 PRBs are available in total. In the downlink, the BS has a high degree of flexibility how it allocates different PRBs to different users, which are denoted by the different colors in Fig. 2.5. In the uplink, the PRBs allocated to one user must be contiguous in the frequency domain. In contrast to Fig. 2.5, which shows a special case discussed in Chapter 4, the allocation in the UL can also vary from TTI to TTI.

In the zoomed view of the downlink in Fig. 2.5, we see that there is a substructure of 12×14 so-called *resource elements* (REs). The beginning of each TTI in the DL (colored gray) is occupied by the *control region*. The *physical downlink control channel* (PDCCH) that conveys the DL and UL scheduling decisions to the mobiles is transmitted there. The remainder of the TTI (colored in blue) is available for payload data transmission and holds the so-called *physical downlink shared channel* (PDSCH). Some of the REs are not available for the PDSCH because they are reserved for so-called *cell-specific reference signals* (CRS), which are transmitted exclusively from one of the BS antenna ports. In the downlink, it is possible to use spatial multiplexing MIMO transmissions as we will explain in more detail in the following sections. In that case, two transport blocks are scheduled and there are two or more *spatial layers* scheduled on the same time and frequency resources.

On a high level, the uplink PHY structure is similar to the downlink with 12×14 available REs as shown in the zoomed view on the right side of Fig. 2.5. The uplink does not have a control region because the scheduling information is signaled from the BS to the MS. The feedback that has to be sent to the BS is multiplexed into the *physical uplink shared channel* (PUSCH) channel. The REs that carry the PUSCH are colored in orange in the zoomed view in Fig. 2.5. The PUSCH is the uplink equivalent to the PDSCH and holds the uplink payload data. If there is no dedicated PUSCH allocation for the MS, the MS uses the so-called *physical uplink control channel* (PUCCH), which occupies a band of a few reserved PRBs on the edges of the spectrum (not shown here), to transmit a reduced amount of feedback information to the BS. Not all REs are available for PUSCH transmission though. For each PRB and TTI with a PUSCH allocation, the MS transmits demodulation reference signals that allow the BS to correctly demodulate the transmission. In addition, the mobiles can be configured to transmit a *sounding reference signal* (SRS) that allows the BS to estimate the frequency-selective uplink channel. If any MS is configured to transmit SRS in a given TTI, all other mobiles must not use the corresponding REs for PUSCH transmission.

In the initial LTE releases 8 and 9, the maximum system frequency bandwidth was 20 MHz. In LTE-Advanced releases 10 and beyond, it is now possible with *carrier aggregation* to bundle multiple not necessarily adjacent frequency bands. With that, a maximum aggregated bandwidth of $5 \cdot 20$ MHz can be achieved, which allows accord-

ingly high per-user peak data rates. In terms of spectral efficiency, which is normalized to the total system bandwidth, carrier aggregation does not make a significant difference. For that reason, we do not consider it further in the following.

2.2.2. The LTE Physical Layer

Data Channel Modulation

The downlink and uplink physical data channels in LTE releases 8 to 11 (i.e., the PDSCH and PUSCH) are transmitted using either QPSK, 16-QAM, or 64-QAM modulation. Figure 2.6 shows the possible complex baseband constellation symbols s in the in-phase and quadrature plane. Based on the number of constellation points in the alphabet, the QPSK, 16-QAM, and 64-QAM modulation schemes allow to convey 2, 4, and 6 bits per symbol, respectively.

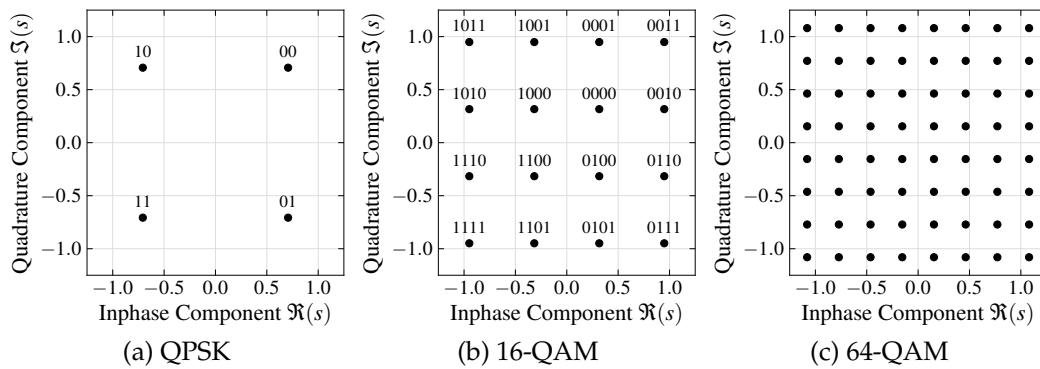


Figure 2.6.: Modulation symbol alphabets as defined in the LTE standard [36.211]

OFDM and OFDMA

Figure 2.7 takes a closer look at the resource grid we already saw in Fig. 2.5. LTE uses *orthogonal frequency-division multiple access* (OFDMA) as its medium access scheme in the downlink and *single-carrier FDMA* (SC-FDMA) for the uplink. Both are based on *orthogonal frequency-division multiplexing* (OFDM), which divides a wide bandwidth into a large number of narrowband *subcarriers*. In LTE, the subcarrier bandwidth is 15 kHz so that 12 subcarriers form a 180 kHz PRB as shown in Fig. 2.5. Having 15 kHz narrow subcarriers entails long symbol durations of $T_U = 1/15\,000\text{ Hz}^{-1} = 66.7\ \mu\text{s}$. This is a key advantage of OFDM systems compared to wideband systems like UMTS. That way, *intersymbol interference* (ISI) resulting from the delayed arrival of previous symbols via longer multipaths is reduced because the length of the channel, that is, the maximum excess multipath propagation delay $\Delta\tau_{\max}$, is usually significantly shorter than the symbol duration. To completely eliminate ISI and to establish orthogonality between the subcarriers, OFDM systems employ a so-called *cyclic prefix* (CP). As shown in

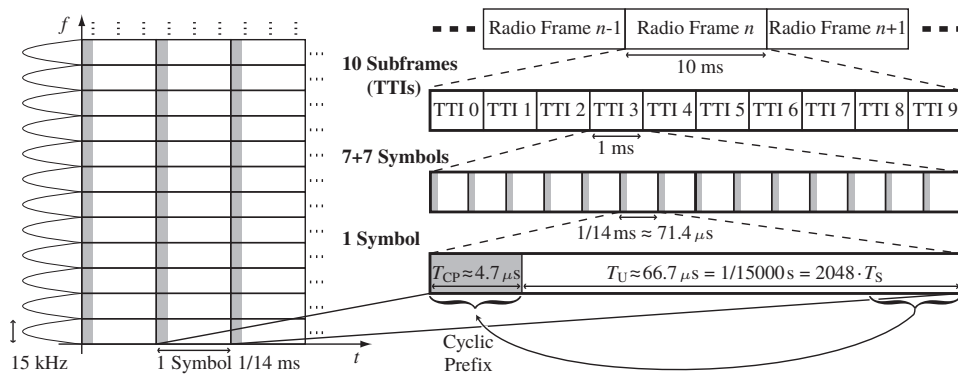


Figure 2.7.: Symbol durations and subcarrier spacing in the LTE resource grid

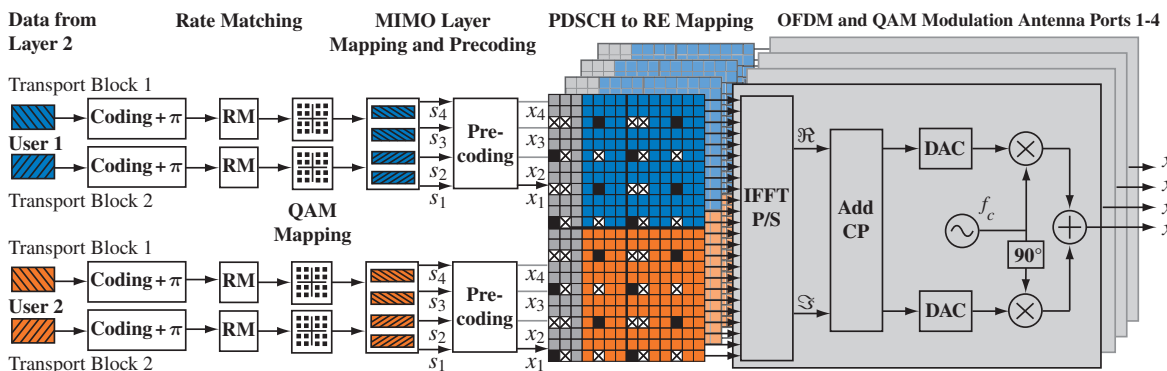


Figure 2.8.: LTE downlink physical layer processing for 2 users on 4 layers

Fig. 2.7, a part of length $T_{CP} > \Delta\tau_{\max}$ from the tail of each symbol is prepended to the symbol before the transmission and at the receiver that part is discarded. Thus, a symbol does not experience ISI from previous symbols, but rather perceives cyclical ISI from itself. This results in orthogonal subcarriers in the frequency domain after performing a DFT [Mol12, Ch. 19] which allows low-complexity broadband receivers [SFFM99]. To keep the overhead in terms of wasted power and transmission time due to the CP as small as possible, LTE defines a normal and an extended CP length. The long CP is only used if the propagation environment causes very long excess delays. Depending on the CP length, there are either 12 or 14 symbols per TTI. Here, and in most practical deployments, we will assume the normal (short) cyclic prefix so that we have 14 symbols per TTI as shown in Fig. 2.7.

Turbo Encoding and HARQ Processing

We will now continue to follow the data flow in the downlink and focus on the processing that happens on the physical layer as shown in Fig. 2.8. When the scheduler in the base station has finished the scheduling for a certain TTI, the transport blocks that have been prepared for all scheduled users are processed by the physical layer according to the resource allocation decisions. LTE follows a *bit-interleaved coded modulation* (BICM) [CTB98] approach that separates the channel encoding from the modulation.

ulator and applies a permutation (denoted by π in Fig. 2.8) to interleave the coded bits. For the channel coding a fixed rate- $1/3$ Turbo coder is used. It operates on so-called *code blocks*, which result from a transport block after segmentation into segments of up to 6144 bits and after attaching a *cyclic redundancy check* (CRC) to the transport block and each segment [36.212]. As we will see in Section 2.2.3, LTE supports a wide range of modulation and coding schemes (MCSs). To support code rates other than $1/3$, LTE employs a so-called *rate matching* mechanism that arranges the $3n$ coded bits (n payload bits called *systematic bits* and $2n$ *parity bits*) in a virtual ring buffer. For a code rate $r > 1/3$ only the corresponding fraction of coded bits in the ring buffer, primarily consisting of systematic bits, is transmitted. For code rates $r < 1/3$ some of the bits in the ring buffer are transmitted multiple times. The starting position in the virtual ring buffer from where bits are selected for transmission, depends on the HARQ *redundancy version* (RV) index. Based on the kind of HARQ retransmission strategy, different RVs can be chosen by the scheduler. With a Chase combining (CC) HARQ strategy, identical bits are retransmitted (i.e., always $RV=0$) for each retransmission. For *incremental redundancy* (IR) retransmissions, different RV indices are used between retransmissions so that with each retransmission different (systematic and) parity bits are selected, hence the name *incremental redundancy*.

QAM and Layer Mapping

The coded bits selected by the rate matcher are bit-level scrambled (not shown in Fig. 2.8) and then mapped to complex baseband modulation symbols (see Fig. 2.6) depending on the modulation scheme chosen by the scheduler. So far, the processing per user has followed one or two parallel flows according to the number of scheduled transport blocks. In the next step, the symbols belonging to transport block 1 and 2 are mapped to MIMO layers. For two layers, the scheduler creates two transport blocks, one for each layer. For more than two layers, transport blocks with two (or three or four) times as many bits are created to be split into equally sized parts that are mapped to the layers. In the example in Fig. 2.8, two users (colored blue and orange) have both been allocated two transport blocks that are mapped to four layers. Taking the modulation order and the number of layers into account, the rate matcher provides exactly as many coded bits as can be mapped to the available resource elements in the PRBs that have been allocated to each user. On each RE, all symbols $\mathbf{s} = (s_1 \ s_2 \ s_3 \ s_4)^T$ belonging to layers 1–4 are transmitted in parallel. However, the symbol vector \mathbf{s} is not transmitted *as is* on the antenna ports. Instead, a precoding matrix \mathbf{P} maps the layer vector \mathbf{s} to the vector of transmit symbols $\mathbf{x} = \mathbf{P}\mathbf{s}$. That way, an arbitrary number of layers $M < N_{Tx}$ can be transmitted and a precoder can be applied based on the available CSIT.

The control channels in the control region at the beginning of the TTI are built in a similar way but with different and typically more robust modulation and coding schemes. Together with the control region symbols, the transmit symbols for all scheduled users as well as the reference symbols, whose REs were not used for PDSCH or control channel symbols, are placed on the time/frequency resource grid. For each of the 14 OFDM symbol durations in the TTI, and separately for each antenna port, the symbols on all

subcarriers are fed into an *inverse fast Fourier transform* (IFFT) to obtain time-domain I/Q symbols. They are serialized and the cyclic prefix is prepended to time-domain symbols belonging to the same OFDM symbol duration. Then, the I/Q symbols are modulated onto the carrier frequency f_c where a 90° phase-shift is applied between the in-phase and quadrature component as shown in Fig. 2.8.

Mobile Station Receiver

As shown in Fig. 2.4, the processing is reversed at the receiver and the recovered user plane as well as RRC data is delivered to higher layers. The processing starts at the radio frequency (RF) transceiver that demodulates the RF signal received on the mobile's antennas. It provides time-domain complex baseband I/Q samples for each antenna port to the baseband processing unit. The baseband processing part of the PHY can be divided into a so-called *inner* and *outer receiver* [SFFM99; MMF97]. Here, we assume that the inner receiver performs functions such as a *fast Fourier transform* (FFT), channel estimation, and MIMO detection to provide *log-likelihood ratio* (LLR) *soft bits* for each transmitted bit to the outer receiver. The outer receiver then decodes the channel code. The inner receiver operates on a resource element level whereas the outer receiver operates on a transport block level. If the Turbo decoder in the outer receiver is able to correctly decode all code blocks belonging to a transport block, the decoded transport block is forwarded further up in the stack and a HARQ *acknowledgement* (ACK) is sent to the BS. If any of the code blocks fails to decode, the whole transport block cannot be decoded and a HARQ *negative acknowledgement* (NACK) is sent to the BS. In this case, the HARQ process stores the LLRs of the failed transport block in order to combine them with future HARQ retransmissions. Not throwing away failed transmissions but keeping the soft bits and thus their information value to combine them with future retransmissions is what distinguishes HARQ from regular ARQ schemes. In Chapter 4 we will show how the performance of the outer receiver is modeled in system-level simulations. In the following, we will focus more on the inner receiver. The exact receiver algorithms are not standardized and can be implementation-specific. Thus, we only describe general concepts here.

By means of an FFT, the inner receiver obtains frequency-domain complex baseband symbols from the I/Q samples provided by the transceiver. As the inner receiver operates on the REs separately, we only consider the vector of received symbols for a certain RE in the following. In accordance with the system model introduced in (2.11), we call the vector holding the symbols received on all antenna ports \mathbf{y} , see also Fig. 2.2. By comparing the received symbols on the CRS positions to the a priori known reference signal sequence, the receiver can estimate the complex baseband channel coefficients. Estimates are available for each antenna pair and by interpolation in time and frequency domain the channel matrix $\hat{\mathbf{H}}$ at an arbitrary RE can be estimated.

Maximum Likelihood Detection

For a SISO link, the transmitted and received symbols per RE as well as the channel coefficient are complex scalars: $y = hs + n$. The task of the detector in the inner receiver is then to determine the transmit symbol $s \in \mathcal{S}$ from the modulation alphabet \mathcal{S} (see also Fig. 2.6) given its knowledge of y and the estimate of the channel h . By assigning an LLR to each corresponding bit, it indicates to the outer receiver the certainty with which each bit has been detected. The optimal detection strategy in the sense that it provides the minimum probability of detection error [CT06] is *maximum likelihood detection* (MLD). In the SISO case, for example, the receiver can undo the effect of the channel (a phase rotation and scaling) on s to obtain an estimate

$$\tilde{s} = h^{-1}y = s + h^{-1}n \quad (2.28)$$

of the transmitted complex baseband symbol. Multiplying by the inverse of the effective channel is often called *zero-forcing* (ZF). Due to channel estimation errors and the noise term $h^{-1}n$, the estimated symbol \tilde{s} will not perfectly match one of the possible constellations. The detector thus has to decide for one of the constellation points. It does so by “rounding” the estimated symbol \tilde{s} to the closest constellation point \hat{s} in the modulation alphabet. In the SISO case, this symbol-by-symbol detection is optimal and MLD is thus very simple [PNG03].

For MIMO transmissions, we also follow the idea to find the closest match between the received signal and the possible transmitted signal. However, the situation is more complicated now because a vector of symbols is transmitted. In this case, the detection decision cannot be performed independently (element-wise) in general. In practice, a transmitter-side SVD-based precoding as shown in (2.19) that achieves orthogonality of spatial layers at the receiver is not possible. The elements of the transmitted symbol vector s of length M will thus interfere with each other. The optimal detection strategy again is MLD [Jal06]. It yields the modulation symbol vector \hat{s} as the best estimate for the transmitted vector s that, after taking the effective channel’s impact into account, is the closest match to the observed symbol vector y :

$$\hat{s} = \arg \min_{s \in \mathcal{S}^M} \|y - \hat{H}Ps\|^2 \quad (2.29)$$

Note that this takes the precoder P applied at the transmitter into account. In LTE codebook-based spatial multiplexing transmission modes, a receiver estimates the MIMO channel \hat{H} using the CRS and is explicitly informed about the applied precoder as discussed in Section 2.2.3. A straight-forward approach to perform MLD would be to exhaustively search over all possible transmit vectors $s \in \mathcal{S}^M$. This is infeasible in practice, as the number of possibilities grows exponentially with the number of layers M , where the size $|\mathcal{S}|$ of the modulation alphabet can be 4, 16, or 64 in LTE. Since MLD has been shown to be NP-hard in the number of layers [Ver89], an efficient MLD algorithm is most likely not feasible. Thus, sub-optimal MLD-like schemes or so-called *linear* receivers [LV89] are used in practice. Jaldén [Jal06] provides an overview of MLD in MIMO channels and analyzes two sub-optimal MLD-like schemes, one of them being the popular *sphere decoding* variant.

Linear Receivers: ZF and MMSE Receive Filters

The linear receiver schemes are called linear because they apply a filter $\mathbf{W} \in \mathbb{C}^{N_{\text{Rx}} \times M}$ to pre-process the received symbol vector \mathbf{y} and obtain a filtered vector as $\tilde{\mathbf{s}} = \mathbf{W}^H \mathbf{y}$, which is a linear operation. From the filtered vector $\tilde{\mathbf{s}}$, the estimate $\hat{\mathbf{s}}$ is then obtained by a component-wise detection process that independently maps each element in $\tilde{\mathbf{s}}$ to the closest constellation point in \mathcal{S} . Thus, an appropriate filter should transform the effective channel so that it becomes diagonal and allows to treat the individual elements of the received vector separately [BCG+07]. One prominent filter choice is again the ZF filter, which for a matrix channel is the *Moore-Penrose pseudo inverse* (denoted by $(\cdot)^+$) of the effective channel $\mathbf{W}^H = (\mathbf{H}\mathbf{P})^+$. It perfectly diagonalizes the channel but also amplifies and colors the noise so that ZF filters are not used in practice. The other prominent linear filter, which is used in various variants in practice, is the linear *minimum mean squared error* (MMSE) filter $\mathbf{W}^H = (\mathbf{P}^H \mathbf{H}^H \mathbf{H} \mathbf{P} + \mathbf{R}_n)^{-1} \mathbf{P}^H \mathbf{H}^H$. It considers both the interference between layers as well as the noise and can be derived as the minimizer of the mean square error [BCG+07]:

$$E\{||\mathbf{s} - \mathbf{W}^H \mathbf{y}||^2\}. \quad (2.30)$$

In the remainder of the thesis, we will only consider linear receivers. On the one hand, there are well-established performance models for our system-level simulation setting. On the other hand, they are still predominant today. For example, MMSE receivers are the baseline assumption in current 3GPP LTE standardization [36.829] as far as mobile station receivers are concerned.

2.2.3. Radio Resource Management in LTE

Radio resource management is an umbrella term that covers many system-level aspects in wireless communications. According to the LTE standard [36.300, Sec. 16], “the purpose of radio resource management (RRM) is to ensure the efficient use [of] the available radio resources and to provide mechanisms that enable [LTE] to meet radio resource related requirements [...]. In particular, RRM in [LTE] provides means to manage (e.g. assign, re-assign and release) radio resources taking into account single and multi-cell aspects.” All of these aspects are under the control of the base station.

In some contexts (e.g., [36.133], [STB11, Ch. 22]) RRM primarily refers to mobility management aspects including measurements and signaling to support cell search and handovers between cells and *radio access technologies* (RATs). Other RRM functionalities listed in [36.300] are *packet scheduling* (which we only touch briefly), *admission control*, and *load balancing* between cells. We do not consider the aforementioned aspects in this thesis and rather focus our discussion of RRM schemes on *dynamic resource allocation* mechanisms. In the following, we will give a brief overview of dynamic resource allocation mechanisms in LTE, concentrating on the downlink. We distinguish between feedback, resource scheduling, and link adaptation aspects for multiple transmission modes. In Chapter 5 we will focus on specific aspects and provide system-level

Table 2.1.: LTE-Advanced DL transmission modes in LTE Rel-11 [36.213, Sec. 7.1]

TM	Description	Max. layers per MS	Reference signal	LTE Rel.
1	Single-antenna transmission	1	CRS	8
2	SFBC transmit diversity	1	CRS	8
3	Open-loop spatial multiplexing	4	CRS	8
4	Closed-loop codebook-based spatial multiplexing	4	CRS	8
5	Codebook-based MU-MIMO mode	1	CRS	8
6	Codebook-based rank-1 precoding	1	CRS	8
7	Rank-1 precoding with user-specific demodulation RS	1	see description	8
8	Non-codebook-based precoding with MU-MIMO support	2	CSI-RS / DM-RS	9
9	Non-codebook-based precoding with MU-MIMO support	8	CSI-RS / DM-RS	10
10	Non-codebook-based precoding for CoMP	8	CSI-RS / DM-RS	11

performance evaluations. In Chapter 6 we will cover *inter-cell interference coordination* methods and *multi-user MIMO* (MU-MIMO) scheduling.

Downlink Transmission Modes

The LTE standard [36.213, Sec. 7.1] in Release 11 defines ten so-called *transmission modes* (TMs) that describe PDSCH transmission, resource allocation, and feedback reporting profiles. Each MS is semi-statically configured to operate in one of the transmission modes. The BS can configure the TM separately for each of the mobiles it is serving. Usually, only a single or very few different transmission modes are employed in one cell in parallel though because the base station often only supports a limited number of modes. We restrict the discussion to downlink transmission modes. On the uplink, which we do not cover in detail here, only one single-antenna transmission mode was available in the initial LTE releases. With Rel-10, a second transmission mode supporting MIMO was introduced. Table 2.1 gives an overview of the downlink transmission modes standardized in LTE up to Release 11.

TM1 is the most basic transmission mode that only employs a single transmission (Tx) antenna at the base station. The MS can use an arbitrary number of antennas for receiving **TM1** transmissions, but with a single Tx antenna, no spatial multiplexing transmissions are possible. All other transmission modes assume the availability of multiple Tx antennas at the BS.

TM2 makes use of two or four BS antennas for *space-frequency block code* (SFBC) transmit diversity operation. SFBC does not allow spatial multiplexing but achieves a diver-

sity gain from using multiple transmit antennas that can be realized without any CSIT. In the absence of CSIT, that is, if the BS cannot adapt its transmissions to the spatial channel characteristics and does not know the rank of the channel, a transmit diversity scheme is a reasonable default transmit strategy. Otherwise, using just one (which?) out of multiple Tx antennas would be an unwarranted optimization assumption in this case. For that reason, SFBC is used for the PDCCH signaling if a BS has multiple Tx antennas. SFBC transmission for the PDSCH is available in all other transmission modes as well by means of a fallback mechanism that can be signaled per TTI, for example, if the BS has no reliable feedback on the spatial channel conditions. We will discuss SFBC in more detail in Section 4.2.3 where we outline the performance modeling and in Section 5.3 where we evaluate its performance.

TM3 and **TM4** are the spatial multiplexing MIMO transmission modes in LTE Rel-8. **TM3** is a so-called *open-loop* scheme whereas **TM4** is a so-called *closed-loop* scheme. In both modes, the mobile station provides a *rank indicator* (RI) feedback. The difference is that in **TM4**, the MS in addition also provides quantized CSI feedback in the form of a *precoding-matrix indicator* (PMI) from a codebook as introduced in the subsequent feedback subsection. For indicated ranks > 1 , it also provides two *channel quality indicator* (CQI) reports for the two spatially multiplexed transport blocks. For **TM3**, only a single CQI and the rank is signaled. The precoder rotates deterministically as shown later when we discuss the performance evaluation in Section 5.5.1.

TM5 and **TM6** both rely on the same precoding codebook and feedback scheme as **TM4** but restrict the transmission to one spatial layer per user. In **TM6**, as in **TM1–TM4**, only one user per PRB is scheduled per cell. Since the rank-1 transmission precoding matrix is just a single column vector, the transmission mode can also be interpreted as codebook-based *beamforming*. In contrast to the previously discussed TMs, **TM5** allows to schedule two users per cell in a MU-MIMO fashion. For that, the BS transmits on identical PRBs using different rank-1 precoding vectors from the codebook with half the usual power.

TM7 allows non-codebook based rank-1 precoding (beamforming) transmissions with user-specific *demodulation reference signals* (DM-RS). As it is not codebook-based, there is no CSI feedback scheme available for *frequency-division duplexing* (FDD) operation so that the use of **TM7** only seems feasible in TDD deployments where channel reciprocity might be exploited to obtain CSIT. We do not consider TDD and **TM7** in this thesis.

TM8, **TM9**, and **TM10** were introduced after the initial LTE Rel-8 standard to allow more antennas and more flexible precoding as well as a dynamic switching between single-user and multi-user MIMO. In contrast to prior transmission modes, they do not rely on CRS for channel estimation and demodulation. Instead, they use new reference signals called *CSI-reference signals* (CSI-RS) and DM-RS. **TM8** is restricted to four antennas, whereas **TM9** and **TM10** allow up to eight Tx antennas and layers per user in SU-MIMO mode or up to four layers in total to a maximum of four users in MU-MIMO mode. **TM10** further supports so-called *coordinated multi-point* (CoMP) transmissions, which we will only briefly discuss in the outlook in Section 6.3.

The chosen transmission mode not only dictates which kind of spatial processing

techniques the BS can employ. It also indicates to the mobile which kind of DL resource allocation signaling and which kind of CSI reporting modes the BS can configure. In the following, we will give only a general description of the feedback mechanism. For the detailed specification of CSI reporting modes, we refer to the standard [36.213, Sec. 7.2].

Feedback

From the discussion in Section 2.1.2, we have seen that the base station needs CSIT to choose suitable precoders. Also, the scheduling and link adaptation algorithms that we will study in the following depend on information regarding the serving channel quality and interference situation. The interference situation and, in an FDD system, also the channel matrix can only be estimated at the desired receiver. To enable channel-dependent resource allocation in the BS scheduler, the MS thus has to transmit feedback information to the BS. As the uplink data rates are limited, finding a suitable trade-off between the level of detail (i.e., quantization, granularity in the time and frequency dimension) and the required signaling bandwidth is an important problem in wireless communication systems [LHL+08]. The general philosophy [LHSH04] underlying the feedback design in LTE is that it is not the actual situation at the MS (e.g., the serving channel matrix and the interference covariance matrix) that is quantized and fed back. Instead, the MS rather tells the BS how it should perform future transmissions by indicating which transmission strategy from a small set of possible choices would be most suitable.

There are three different CSI feedback components in LTE: the CQI, the PMI, and the RI. In all transmission modes the mobile provides a *channel quality indicator* (CQI) which is a 4-bit value as shown in Table B.1 in the Appendix. The CQI indicates the highest-rate modulation and coding scheme that the mobile would be able to receive with a *block error ratio* (BLER) of at most 10% given the current serving channel and interference condition [36.213, Sec. 7.2.3]. In the end, the CQI corresponds to a certain SINR value. But by not feeding back a quantized SINR (or even raw channel matrices), the MS vendors can differentiate by implementing more advanced receivers that might support higher rates given a certain SINR. This is an application of the philosophy outlined above. The CQI feedback can be *wideband* (i.e., only one value for all PRBs) or *subband-specific* to allow the scheduler to exploit the frequency-selective channel or interference situation. With subband-specific feedback, multiple PRBs are aggregated into a so-called *subband*. Only one CQI is provided per subband reducing the signaling overhead. In addition, the subband CQI is differentially encoded with 2 bits per subband relative to the 4-bit wideband CQI.

The MS does not need to have an ongoing PDSCH transmission to be able to estimate the channel for feedback purposes because it can use the cell-specific reference signals to estimate the serving MIMO channel matrix \hat{H} . That way, it can also estimate the inter-cell interference because the CRS are interfered by other-cell PDSCH transmission as neighboring cells are usually configured with shifted CRS patterns avoiding CRS collisions in synchronized systems. Depending on the transmission mode, the MS

knows what transmitter-side spatial processing the BS can perform. It can thus apply the same processing to the channel \hat{H} and then evaluate how well it would be able to receive a corresponding transmission with its receiver implementation. As there is no actual transmission, it cannot directly use its standard inner receiver implementation but has to resort to link performance modeling techniques similar to those that we will present in Chapter 4 for system-level simulations.

The *rank indicator* (RI) is provided in spatial multiplexing transmission modes. In order to reduce signaling overhead, the LTE standard only allows to provide one RI value for the whole system spectrum (“wideband”). This is justified because the rank of the channel matrix depends on the richness of the scattering environment, which is not very frequency-selective. It is also expected to change only slowly (see also the discussion in Sections 2.1.1 and 2.1.2) so that rank updates are typically less frequent than CQI or PMI updates. Note that despite a numerically full-rank MIMO channel matrix, a mobile might want to signal a low-rank RI because the weaker eigenmodes are not strong enough or not well-aligned with the available precoders.

For the closed-loop spatial multiplexing in TM4, but also for the codebook-based rank-1 precoding in TM5 and TM6, the MS provides a so-called *precoding matrix indicator* (PMI). As the BS has no direct knowledge of the channel matrix, it is the mobile’s task to pick a suitable precoder from a pre-defined codebook of precoding matrices and signal it in form of a PMI to the BS. Taking into account the rank as indicated by the RI, the MS picks that precoding matrix from the codebook that yields the highest expected data rate taking all layers together. Depending on the transmission mode and the feedback reporting configuration, the PMI can be signaled in a wideband or subband-selective way. If the MS indicates a rank higher than one, it also has to provide two potentially different CQI values (except for TM3): one for each transport block considering the layer mapping according to the signaled rank.

Since the MS can estimate the pure MIMO channel \hat{H} , the MS can compute all candidate effective channels $\hat{H}P(\text{PMI}, \text{RI})$ as well as its hypothetical receive filters W given the current interference situation. In a way, the PMI represents a very coarse quantization of the eigenvectors of the channel (this is especially obvious for a rank-1 PMI because with a single-layer transmission, the best strategy is to beamform into the dominant eigenmode direction of the channel as discussed in Section 2.1.2). Table B.7 in the Appendix lists all four possible rank-1 and rank-2 precoding matrices for the case with two Tx antennas at the BS. In Tables B.3 to B.6, the 16 PMIs available for rank 1–4 transmission in the case of four Tx antennas at the BS are shown. We will discuss further properties of the LTE MIMO codebook in Chapters 5 and 6.

An extension of the codebook-based MIMO feedback scheme is applied to TM9 and TM10. The channel estimation is not performed on the CRS but on special CSI reference signals called CSI-RS which allow to estimate the channel to up to eight Tx antennas with less overhead compared to the legacy CRS. To support eight layers, the PMI scheme has been extended in Rel-10. We will not discuss these extensions in this thesis. While the feedback in these transmission modes and also in TM8 is restricted to the codebook, the BS can apply arbitrary precoders in the downlink, which is especially

useful to support MU-MIMO operation in a seamless way. This is accomplished by providing so-called user-specific DM-RS that are precoded in the same way as the PDSCH data transmission so that the MS can use them directly to estimate the effective channel needed for successful demodulation and MIMO detection.

The feedback is reported on the uplink either in a *periodic* fashion as often as every 2 ms for the CQI or upon explicit request by the BS in an *aperiodic* manner. The exact reporting configuration options are TM-specific and quite complex. For the simulations later in this thesis, we will indicate whether we consider wideband or subband CQI and PMI feedback and how often we report CQI, PMI, and RI values.

Scheduling

The term *scheduling* in the LTE context refers to *user scheduling*, *packet scheduling*, and *resource scheduling*. All of these operations are performed jointly each 1 ms TTI by the base station. According to the packets the BS holds in its queues for each user, and depending on the need to perform HARQ retransmissions, the BS decides which users it wants to schedule in the current TTI. Taking into account the QoS requirements associated with the packets waiting for a certain user, the BS decides which of the packets are multiplexed onto the physical resources and how many bits are needed in total. The feedback provided by the users gives an indication of their individual channel conditions. With that information, the BS determines how many PRBs and layers it allocates to each user. Most important, frequency-selective feedback allows to give specific parts of the frequency spectrum to those users who would benefit most. Also, users with temporarily bad channels might be postponed to subsequent TTIs if their QoS requirements allow. If, in addition, precoding feedback is available, the BS can choose suitable precoders and, in the case of MU-MIMO operation, it can select groups of users that are co-scheduled onto the same PRBs.

From that overview, it becomes clear that the BS has to solve multiple scheduling problems at once. A further complication to real-world scheduling in LTE arises from restrictions on the available *downlink control information* (DCI) signaling. In an attempt to keep the required per-user signaling bandwidth on the PDCCH low, only certain resource allocation choices can be signaled depending on the transmission mode [36.212, Sec.5.3.3]. Satisfying all of these constraints at once, and deciding for the best strategy according to some metric on a 1 ms time scale is an extremely difficult problem (some variants of which have explicitly been proven as NP-hard [LCK+09; ZPR12]). For that reason, heuristic approaches are used both for evaluations as well as in practice. For the evaluations in this thesis, we implemented a few common heuristics, which we will explain in the following. To simplify the implementation, we do not model the constraints arising from the signaling limitations, and we will not discuss them in detail here.

Fig. 2.9 shows an idealized situation during one TTI where the SINR for three users, as measured at the MS, is plotted over the frequency bandwidth spanning 25 PRBs. Note that the actual information available at the BS from the CQI feedback is coarser. It is

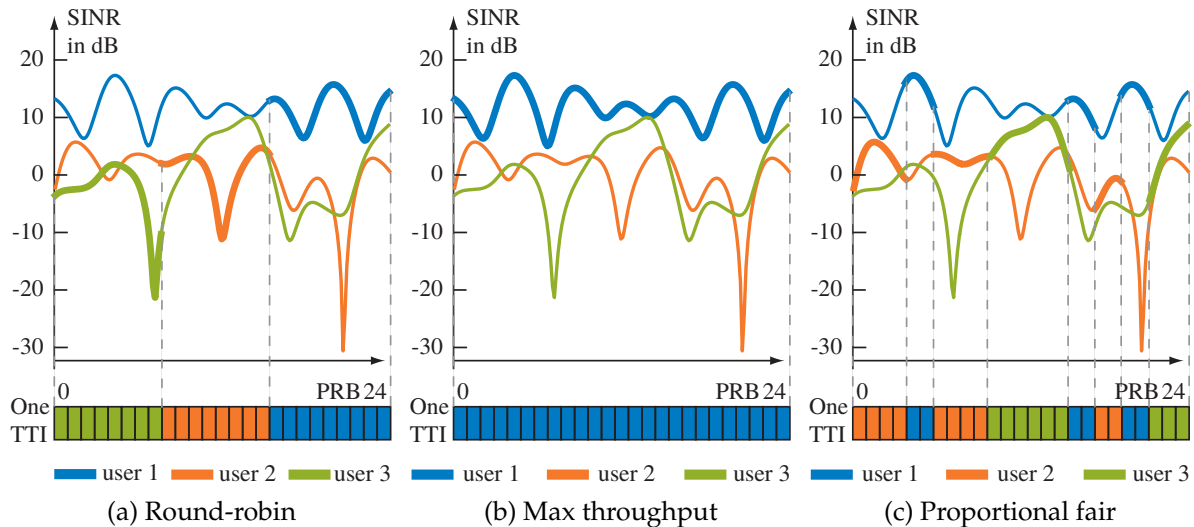


Figure 2.9.: Different scheduling strategies in highly frequency-selective conditions

also not instantaneously available because of an inherent delay due to the signaling periodicity and the feedback processing time. Three different generic strategies are shown in Fig. 2.9. The first strategy, shown in Fig. 2.9a, is called *round-robin* (RR). It is not channel-aware and simply allocates the physical resources in a round-robin fashion to the users. The selected PRBs are shown at the bottom with different colors representing different users. In the example, the users are scheduled all at the same time (i.e., in the same TTI) in equally-sized blocks in the frequency domain. Variants of RR could alternate the users in the time-domain only (i.e., allocating all resources to one user per TTI). The fading channel curves are shown with thick lines for the part of the spectrum where a user is scheduled. RR is *resource-fair* as each user gets an equal share of resources. It is, however, not necessarily *throughput-fair* since users with good channel conditions, like user 1 in the example, will get higher rates if the link adaptation correctly assigns higher rate MCSs to users with good channel conditions. The highest cell throughput, but also the most unfair resource allocation, is obtained by a method called *maximum throughput* scheduling that gives each PRB to the user who experiences the best channel conditions on that PRB. This can be extremely unfair, as shown in Fig. 2.9b, if a single user, who might be very close to the BS, always has a better channel than all other users. Under max throughput scheduling it would get exclusive access while all other users are not served at all.

A compromise between maximum throughput and throughput-fair scheduling is *proportional fair* (PF) scheduling [Hol00; Hol01] as depicted in Fig. 2.9c. Originally proposed by Tse in 1999 [Tse99], it has become a standard metric for channel-dependent scheduling. It selects that user k^* from the set \mathcal{K} of all users for transmission on a PRB with index f_{idx} who offers the maximum ratio of the expected rate on the PRB divided by the user's past overall throughput $\bar{\gamma}_k$:

$$k^*(f_{\text{idx}}) = \arg \max_{k \in \mathcal{K}} \frac{\log [1 + \text{SINR}(k, f_{\text{idx}})]}{\bar{\gamma}_k}. \quad (2.31)$$

That way, users are scheduled on those PRBs where they promise to have good conditions, while at the same time users with low rates in the past still receive a fair share of resources. Figure 2.9c shows how the scheduler can obtain a *multi-user diversity gain* [KH95] by “riding on the peaks” [Tse01] of the fading envelopes of the individual users’ channels. In Section 5.1.1 we will discuss further aspects of the PF scheme and present performance evaluations.

The exact scheduling algorithm is not standardized in LTE allowing for differentiation by BS vendors. Capozzi et al. [CPG+13] present a current survey of downlink scheduling schemes for LTE considering also QoS implications. In the uplink, the BS also controls the resource allocation and signals its decisions to the mobiles who then multiplex packets from potentially different QoS queues to the indicated resources. There are, however, a few differences in the uplink. First of all, only contiguous PRB blocks can be allocated so that none of the allocations shown in Fig. 2.9c would be a valid UL allocation. Second, the BS as the receiver in the UL can estimate the channel itself but needs SRS transmissions (see Fig. 2.5) from the mobiles if there are no data transmissions that allow to estimate the channel. Third, the mobiles have to indicate their need for uplink allocations to the BS by means of scheduling requests and buffer status reports indicating the amount of data they want to transmit. In addition, they provide power headroom reports to inform the BS about their power budget for PUSCH transmission.

Link Adaptation and HARQ Retransmission

An important task that goes hand-in-hand with the scheduling is *link adaptation*. Based on the amount of backlogged bits in the queues, the scheduler allocates a certain number of PRBs to a user. The amount of payload bits that can be transmitted on these PRBs, however, also depends on the number of layers and the MCS. Choosing the MCS, that is, the modulation and the code rate, is called link adaptation. The higher the modulation order and the higher the code rate, the larger is the possible data rate, but the more susceptible the transport block is to corruption on the channel resulting in a block error. With the help of the CQI feedback, the link adaptation thus aims to select the highest MCS that it deems to be robust enough given the estimated channel conditions. The mapping from CQI to MCS is not straightforward though. The MCS applies to one transport block as a whole, that is, it is constant over all PRBs but potentially different between two transport blocks transmitted with spatial multiplexing. The MCS index can be chosen from 29 possible values whereas the CQI only provides a granularity of 15 values. In addition, for the PRBs to which a transport block is allocated, the mobile might have indicated different subband-specific CQI values. The link adaptation algorithm in the BS thus might have to aggregate multiple CQIs for which it can use link performance models similar to those presented for the computation of effective SINRs in Chapter 4. The BS typically does not aim at a BLER close to zero because if the link adaptation is too cautious, the corresponding data rates are too low. Since a too-aggressive MCS selection can usually be corrected by a HARQ retransmission, the BS trades-off a certain block error probability against higher possible rates. A value of

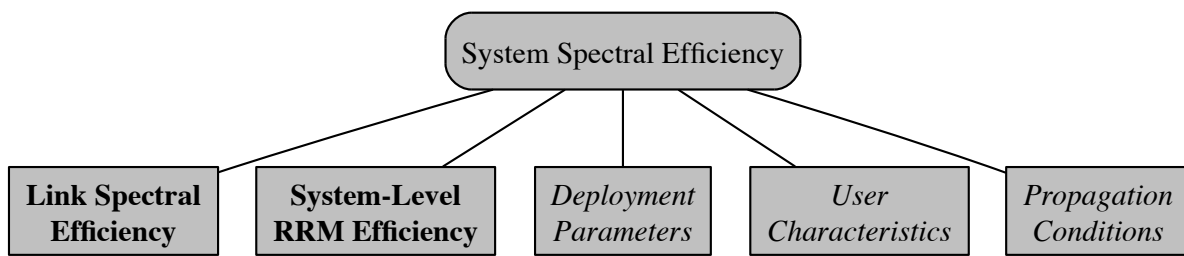


Figure 2.10.: High-level overview of determinants affecting system spectral efficiency

10 % is often considered to be a good BLER target and thus is also the assumption that the MS has to use for the CQI feedback [36.213, Sec. 7.2.3]. Since the BS only sees the reported CQI values but neither knows the actual channel nor the exact performance characteristics of the receiver, it usually uses the HARQ feedback as an additional input to achieve the desired target. The respective mechanism is called *outer-loop link adaptation* (OLLA) and will be discussed in more detail in Chapter 5 together with a performance evaluation of link adaptation.

In general, either adaptive modulation (and coding), as in the LTE downlink, or power control with a fixed modulation can be used to achieve a high link spectral efficiency over a fading channel [CG01]. Adapting the MCS and using a fixed power allows to operate the power amplifiers more efficiently [DPS11] and also leads to more predictable interference levels in the system. Note that instead of link adaptation, a fixed MCS could be employed. Then HARQ retransmissions, which allow the receiver to obtain CC gains or to receive additional parity bits (IR mode), could be used as a correction mechanism. In fact, HARQ can be seen as an *ex post* link adaptation, that is, an adaptation that occurs only in case it turned out to be necessary, compared to the *ex ante* link adaptation mechanism that is based on a prediction. However, predictive link adaptation allows for shorter delays and the needed CQI feedback is an important enabler for achieving multi-user diversity gains.

2.3. LTE Link Spectral Efficiency

After introducing general bounds on channel capacity and implementation details of LTE in the first two sections of this chapter, we now begin to focus our attention on the spectral efficiency achievable in LTE. Figure 2.10 shows a modified version of the overview already presented in Fig. 1.4 where we left out the dependency on the spectrum and the number of sites. Instead, we focus on the spectral efficiency, which is independent of these two aspects. Even before introducing *deployment parameters*, *user characteristics*, and *propagation conditions* in Chapter 3, and without resorting to system-level simulations, we can outline a number of reasons why link spectral efficiencies in LTE inherently fall short of channel capacities.

For a moment, we will only consider a single link. In the remainder of the thesis we will then always consider the system-level performance, which also crucially depends

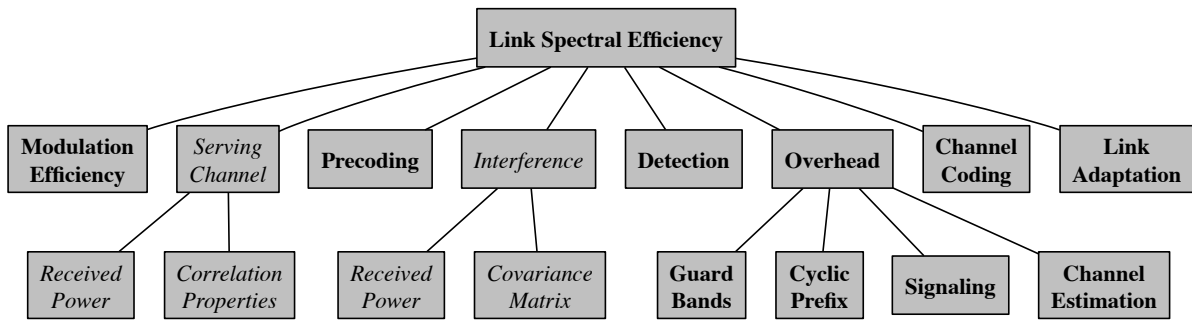


Figure 2.11.: Selected determinants affecting link spectral efficiency

on the performance of the chosen RRM schemes on a system-level. Figure 2.11 gives an overview of the main factors that determine the link spectral efficiency in LTE. We consider the serving channel and the interference as given here, but the overhead will be discussed separately. The remaining factors (modulation, channel coding, link adaptation, precoding, and detection) are a testament to the fact that we are not only interested in channel capacities but in transceivers that can transmit data over actual channels. In the following, we will outline what causes the gap between channel capacity and achievable rates. Most of the reasons are implementation-specific. To illustrate the size of the individual losses, we present a quantitative example that we derived from results published by Schwarz et al. [SSR11; RCMW11]. An early comparison between the LTE link efficiency and capacity can be found in [MNK+07]. A more recent one is presented in [LPT+12], and the authors in [MCR11; LR13] further provide measurement results from the field.

2.3.1. Suboptimal Transmitters and Receivers

In Section 2.1.2 we introduced the MIMO channel capacity as the upper bound on the rate with which we can transmit over a given channel. Among the necessary assumptions were Gaussian-distributed channel inputs and perfect CSIT. The latter allows performing an SVD-based precoding (2.21) with waterfilling in the spatial domain and, now that we consider multiple frequency subbands, also in the frequency domain. Neither of these preconditions holds for the transmitter side of an actual system such as LTE. Furthermore, also the receiver does not perform optimally.

Table 2.2 shows the losses between different levels of throughput bounds as derived from a publication by Schwarz et al. [RCMW11]. There are four link scenarios representing a 1×1 SISO link, a 2×2 MIMO link with *open-loop spatial multiplexing* (OLSM) and a 4×4 MIMO link with *closed-loop spatial multiplexing* (CLSM), as well as an 8×8 MIMO link again with CLSM. Row a) indicates the net capacity, that is the channel capacity according to (2.21) from which a CRS overhead depending on N_{Tx} is subtracted. Otherwise, the net capacity grows linearly with the number of possible layers since the authors assume uncorrelated MIMO channels. Compared to the net capacity rates, the first loss in row b) results from a restriction of the possible precoders to the LTE

Table 2.2.: Exemplary transceiver performance with net rate and attribution of individual losses compared to capacity at 10.4 dB according to [RCMW11]

Measure	1x1		2x2 OLSM		4x4 CLSM		8x8 CLSM	
	Rate	Loss	Rate	Loss	Rate	Loss	Rate	Loss
a) Net capacity	2.75	–	5.25	–	10.00	–	19.00	–
b) Codebook precoding	2.71	1.5%	5.10	2.9%	9.50	5.0%	16.00	15.8%
c) Linear (ZF) receiver	2.71	0%	4.25	16.2%	8.70	8.0%	13.40	13.7%
d) Subband precoding	–	–	–	–	8.20	5.0%	–	–
e) BICM 4/16/64-QAM	2.52	6.9%	4.00	4.8%	7.60	6.0%	12.50	4.7%
f) Finite BL, discrete CRs	2.20	11.6%	3.45	10.5%	6.75	8.5%	11.00	7.9%
g) Turbo coding (est.)	1.95	9.1%	3.10	6.7%	6.00	7.5%	9.90	5.8%
Total loss / overhead	0.80	29.1%	2.15	41.0%	4.0	40.0%	9.10	47.9%

codebook due to the quantized PMI-based feedback in LTE. This includes the inability to perform waterfilling in the spatial domain as well as in the frequency domain, which is the reason for the difference observed in the SISO scenario. The next performance loss in row c) results from the use of non-MLD receivers in case of spatial multiplexing. As described in Section 2.2.2, linear receivers (like the assumed ZF filter in this case) are only optimal for single-layer transmissions. In a spatial multiplexing setting, a part of the mutual information rate available after b) is lost without an MLD receiver. Note from the table that the loss is highest for open-loop spatial multiplexing because even the limited number of precoders in the LTE codebook achieves a certain degree of orthogonalization between the spatial layers. The authors of [LPT+12] further observe that the loss from linear receivers grows with a higher number of antennas and with the degree of correlation in the MIMO channel for high rank. The next loss shown in row d) results from a limited precoding granularity assumed for the 4×4 CLSM example, where the same precoder has to be applied to all PRBs in one subband (here of size 6 PRBs). BICM with fixed modulation alphabets, rather than Gaussian signaling, results in an additional loss compared to capacity. The results in row e) are based on bounds for BICM provided by Caire et al. [CTB98]. The channel coding contributes losses as seen in row f) due to limited code block sizes and because only certain code rates are available with the 29 MCS choices in LTE. Finally, row g) summarizes the remaining gap to link-level simulation results, which the authors attribute to the loss resulting from the Turbo coding performance [RCMW11].

Overall, we observe that in the examples between 29% and 48% of the theoretically-available channel capacity cannot be realized due to implementation losses resulting from the various link-level design choices. What is not even considered here, are losses due to channel estimation errors and RF impairments at the transmitter and the receiver such as the *error vector magnitude* (EVM) [36.101; 36.104] defining the allowed deviations from the ideal transmit modulation symbols. What is also not yet accounted for, are losses caused by the signaling and channel estimation overhead, which we will consider in the following.

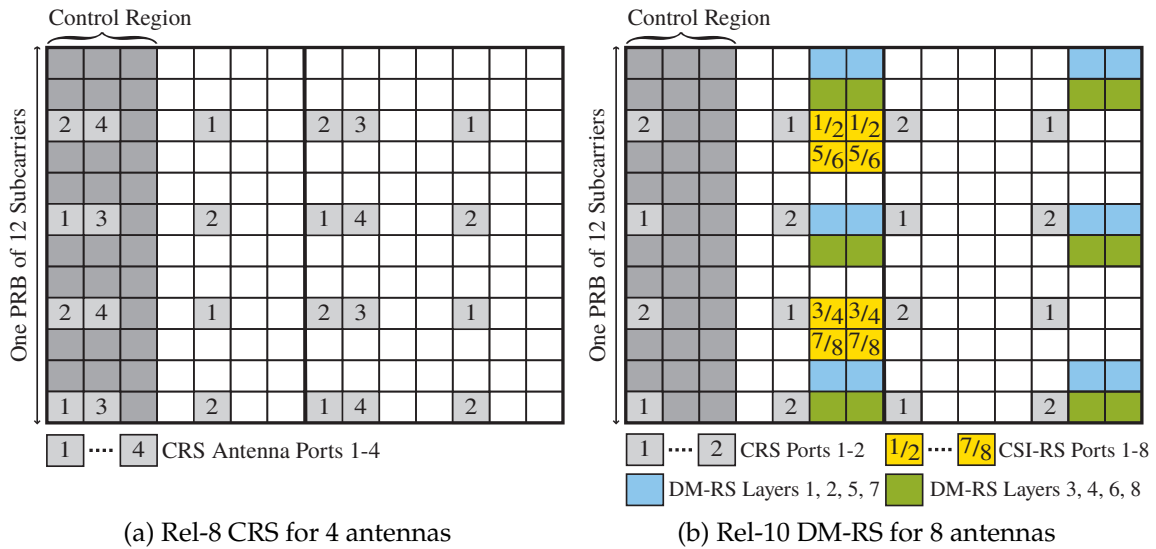


Figure 2.12.: Comparison of exemplary reference signals in LTE Rel-8 and Rel-10

2.3.2. Overhead in the LTE Physical Resource Grid

Figure 2.12 shows a detailed picture of the downlink control signaling and reference signal overhead for the smallest schedulable unit in LTE, that is, one PRB in the frequency and one TTI in the time domain. Only the unshaded resource elements in the 12×14 elements resource grid are available to carry the PDSCH data transmission. In Fig. 2.12a, which illustrates a similar configuration as the one already shown for LTE Rel-8 in Fig. 2.5, the first three OFDM symbols (columns) are occupied by the *control region*. The control region can be between one and four symbols long (3 symbols in the example in Fig. 2.12) and carries the PDCCH indicating the DL and UL scheduling decisions. Also, HARQ feedback for the uplink (via the *physical hybrid-ARQ indicator channel* (PHICH)) and an indicator for the control region length (via the *physical control format indicator channel* (PCFICH)), which can change each TTI, are transmitted there. The REs shaded in light gray are used for CRS, which are transmitted from the indicated antenna ports while the other antennas are muted. That way, the MS can estimate all entries in the MIMO channel matrix. Depending on the BS antenna array, CRS REs for only one antenna port, two antenna ports, or all four antenna ports can be present. Note that, in an attempt to limit the additional overhead when four Tx antennas are used, less CRS REs are provided for antenna ports 3 and 4 compared to the ports 1 and 2. Still, the REs reserved for the control region and for the CRS occupy a significant part of the resource grid. Table 2.3 lists the number of occupied REs for different configurations. We provide the total number of occupied REs as well as the relative overhead compared to the 12×14 REs available per PRB and TTI. Also included in the table, but not shown in Fig. 2.12, are overhead contributions (computed on a pro rata basis relative to a 10 MHz bandwidth) for the primary and secondary synchronization signals (PSS and SSS) and broadcast (PBCH) channels that are transmitted on the six center PRBs with a periodicity of 5 and 10 ms, respectively.

Table 2.3.: PHY system design overhead for different number of antenna ports (with 12×14 REs, normal cyclic prefix, 10 MHz system bandwidth)

Resource element overhead	1 Tx ports Rel-8		2 Tx ports Rel-8		4 Tx ports Rel-8		8 Tx ports Rel-10	
	L=1	L=3	L=1	L=3	L=1	L=3	L=1	L=3
Control region (L={1,3})	10	34	8	32	8	28	8	32
CRS	8		16		24		16	
DM-RS	0		0		0		24	
CSI-RS (5 ms period)	0		0		0		0.2 · 8	
Sync and PBCH	1.1		1.1		1.1		1.1	
Total overhead in REs	19.1	43.1	25.1	49.1	33.1	53.1	50.7	74.7
Rel. overhead	11.3 %	25.6 %	14.9 %	29.2 %	19.7 %	31.6 %	30.2 %	44.4 %

In Figure 2.12b and in the last column of Table 2.3, the overhead for a Rel-10 configuration with eight Tx antennas is illustrated. Since the Rel-8 CRS scheme would not scale with a reasonable overhead to support up to eight antennas, a new hybrid reference signal scheme has been introduced for Rel-9 and beyond. To support legacy mobiles, the CRS have to be kept, but typically only two CRS antenna ports are configured. There are now two separate reference signals to support the two main purposes. First, reference signals are needed to estimate the channel for feedback generation. This is accomplished with low-overhead CSI-RS, which are only transmitted every 5–80 ms. The reference symbols from two antenna ports are transmitted on the same pair of REs separated by a so-called *orthogonal cover code* (OCC). The same principle is applied to the second class of reference signals called DM-RS, which are used as demodulation reference signals. Since channel estimation is now possible with the CSI-RS, the DM-RS can be precoded in the same way as the PDSCH REs allowing the BS to use arbitrary precoders. Also, the amount of occupied REs now scales with the number of actually used layers and not with the number of Tx antennas. As shown in Fig. 2.12b, 12 REs are needed for a two-layer transmission and 24 REs for 3–8 layers.

2.3.3. Total PHY Loss and Link Spectral Efficiency

The bottom line of Table 2.3 reveals that the relative overhead in the LTE resource grid varies between 11 % and 44 % in the different example configurations. Within the remaining resources, data transmissions can take place, which are subject to the link-level losses exemplified in Table 2.2. As big as they are, these two classes of losses are not the only PHY-level factors reducing capacity. The normal cyclic prefix shown in Fig. 2.7 reduces the useful symbol duration by a factor of approximately 6.67 % on average. In the frequency domain, *guard bands* around the OFDM subcarriers protect adjacent bands. For example, in a 10 MHz LTE deployment only 50×180 kHz PRBs are available occupying 9 MHz of spectrum which contributes another 10 % of overhead.

All these components are multiplicative, so that the total PHY loss can be computed, for example, in the case of a Rel-8 4×4 CLSM configuration with a control region length of $L = 3$ as:

$$\text{Total PHY Loss} = 1 - \underbrace{(1 - 0.4)}_{\text{Transceiver}} \cdot \underbrace{(1 - 0.316)}_{\text{RE Overhead}} \cdot \underbrace{(1 - 0.1)}_{\text{Guard Bands}} \cdot \underbrace{(1 - 0.067)}_{\text{Cyclic Prefix}} = 65.5\%. \quad (2.32)$$

Apparently, from the channel capacity in (2.21) more than 65% can be lost on the physical layer alone, and that is even assuming that RRM functions, for example, the precoding feedback and link adaptation, work optimally. On the way to the application layer, more overhead is accumulated in higher layers of the protocol stack. Beginning with CRC checksums for each codeword, headers for MAC, RLC, and PDCP packet data units (PDUs), as well as compressed IP headers reduce the net data rates available to applications. The relative overheads depend heavily on packet sizes and bearer configurations so that we will not discuss them in detail here. Note that we consider CRC checksums and basic LTE headers as overhead in our simulations even though they typically do not contribute significantly to the total overhead.

2.4. Evaluating System Spectral Efficiency by System-Level Simulation

Despite the sizable losses discussed in the previous section, an actual LTE mobile station can—*in theory* (and under laboratory conditions)—still achieve very high peak data rates. The technical specification (TS) [36.306] defines several device categories with different peak rate requirements. For example, in a 20 MHz bandwidth, an LTE category 3 mobile needs to support a DL data rate of 100 Mbit/s, while a category 5 mobile has to be capable of receiving 300 Mbit/s in the DL with the same frequency bandwidth. These peak rate examples correspond to link spectral efficiencies of 5 bit/s/Hz and 15 bit/s/Hz, respectively. Such peak data rate categories are important marketing figures for device manufacturers, but they do not mirror the data rates that can be *realistically* reached *on average* in a multi-cell and multi-user system. Many propagation channels do not even offer such high channel capacities due to low SINRs or correlated MIMO channels. It is thus important to consider realistic propagation conditions. But also the RRM schemes play an important role because SINRs and link conditions are heavily influenced by RRM functions like scheduling and precoding. In this section, we will thus motivate why we need to resort to *system-level simulations* to assess the system spectral efficiency of an LTE system.

Figure 2.10 shows that besides the link spectral efficiency and the BS deployment and user characteristics, the *system-level RRM efficiency* is an important determinant for the overall system spectral efficiency. In contrast to the link spectral efficiency, it is impossible to separate the efficiency impacts of individual RRM methods because they are closely intertwined. This makes it hard to use analytical models and motivates the simulation-based evaluation of RRM schemes. In addition, the RRM performance,

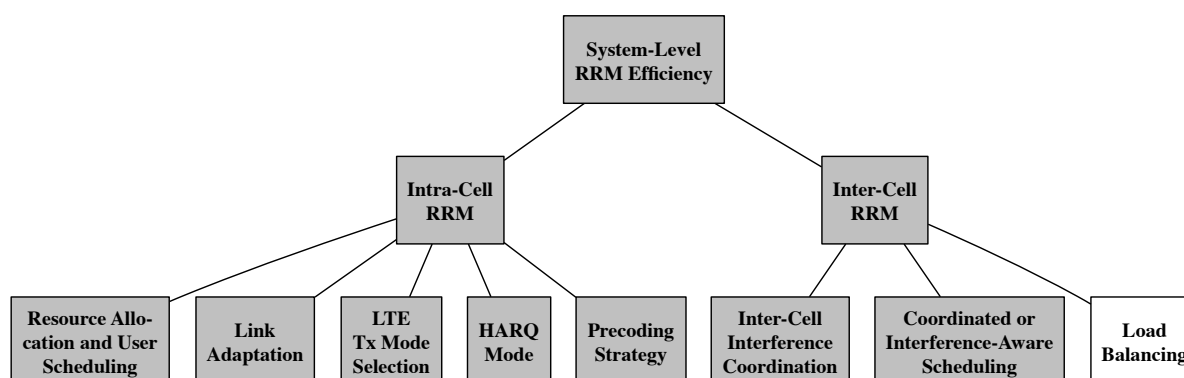


Figure 2.13.: Selected intra-cell and inter-cell RRM aspects affecting spectral efficiency

and thus the spectral efficiency of the whole system, depends on the propagation conditions on all links. To provide a realistic mix of propagation conditions, we make use of sophisticated channel models, which we will discuss in detail in Chapter 3.

2.4.1. System-Level RRM Spectral Efficiency Determinants in LTE

A functional cellular communication system is more than just a conglomerate of individual links. It needs complex RRM mechanisms to operate each link—but especially the system as a whole—in an efficient way. Figure 2.13 gives a coarse classification of selected RRM methods that affect the system-level RRM efficiency. We distinguish between methods that are, on the one hand, primarily focused at the users of the cell served by one base station (*intra-cell RRM*) and, on the other hand, methods that are meant to manage resources among multiple cells (*inter-cell RRM*).

Already at this stage, we observe that the dependencies between the different classes of RRM schemes are not as clear-cut as Fig. 2.13 suggests: Apart from inter-cell load-balancing and handover aspects, which we will not consider further in this thesis, the main focus of inter-cell RRM is to manage inter-cell interference. But the inter-cell interference also influences the intra-cell RRM decisions in the considered cell and vice versa. Also, there is a strong interdependence within the intra-cell RRM methods. The selection of the LTE transmission mode for each user influences the possible resource allocations and precoding schemes that can be used for that user. It also determines which link adaptation feedback is available and that in turn dictates what a scheduler knows about the channel. On the other hand, the interference resulting from the scheduling and precoding strategies in one cell influence the reliability of the link adaptation in neighboring cells, which, as a result, alters the scheduling in those cells. From these examples, it should be obvious that the performance of one RRM scheme can only be judged together with other RRM functionalities in the context of a multi-cell system, and considering channel conditions. Besides, the performance of a conceptually simple RRM operation like, for example, zero-forcing precoding, crucially depends on the accuracy and timeliness of feedback signaling. Other algorithms, for example, for finding optimal sets of co-scheduled users, are computationally so

demanding that heuristics have to be used in a real-time environment. Consequently, simpler models that do not consider these real-world limitations, can only provide upper bounds but no reliable estimate of real-world performance.

All of these factors and interdependencies make it impossible to provide a meaningful *analytical* system-level performance evaluation of RRM schemes. We will thus use a holistic approach that evaluates all RRM schemes jointly. Simulations allow to provide realistic inputs to concrete implementations of RRM schemes, which can even emulate the algorithms that are used in the field. The importance of good simulations in the face of an ever-increasing complexity in wireless systems has also been acknowledged by academia [DHL+11]. We emphasize that, in order to obtain reliable and comparable simulation results, careful modeling and adherence to generally-accepted simulation methodologies is important. In the following, we will introduce the ITU's IMT-Advanced simulation guidelines, which are considered to be the current gold standard of system-level simulation.

2.4.2. System-Level Evaluation Methodology

Simulations of wireless systems like LTE are essential in the research and standardization phase when the technology is still being developed. At this stage, prototypes are often not available and testing different candidate features in the field would be too expensive and time consuming. In comparison, simulations can be set up easily and deliver quick results. For the simulation results to be meaningful and reliable, however, both the channel, as well as the investigated wireless system, have to be modeled accurately enough. The accuracy of the channel model is important because the performance gains over legacy systems delivered by LTE systems mainly stem from a better exploitation of the radio channel's selectivity in the time, frequency, and spatial domain (i.e., the fast-fading MIMO channel). Consequently, the ITU adopted a well-established spatial channel model developed within the Wireless World Initiative New Radio (WINNER) project when it defined detailed guidelines [M.2135] for the evaluation of LTE-Advanced and IEEE 802.16m as IMT-Advanced compliant technologies.

In its "Guidelines for evaluation of radio interface technologies for IMT-Advanced", first published in 2008 in report ITU-R [M.2135s; M.2135], the ITU defined the general evaluation methodology, the relevant figures of merit, test environments, scenarios, and corresponding channel models, which we will present in Chapter 3. The aim of that report was to define test environments that are representative of expected use cases for IMT-Advanced systems and to allow comparable evaluation results by providing well-defined evaluation guidelines. In this thesis, we will focus on system-level simulations in accordance with report [M.2135] because the ITU requires using simulations for evaluating the cell spectral efficiency. Other (simpler) evaluation methods can be used to assess characteristics like peak rates or latency bounds. In essence, system-level simulations consider the performance across multiple layers of the protocol stack for a whole cellular system consisting of many communication links in multiple cells. In

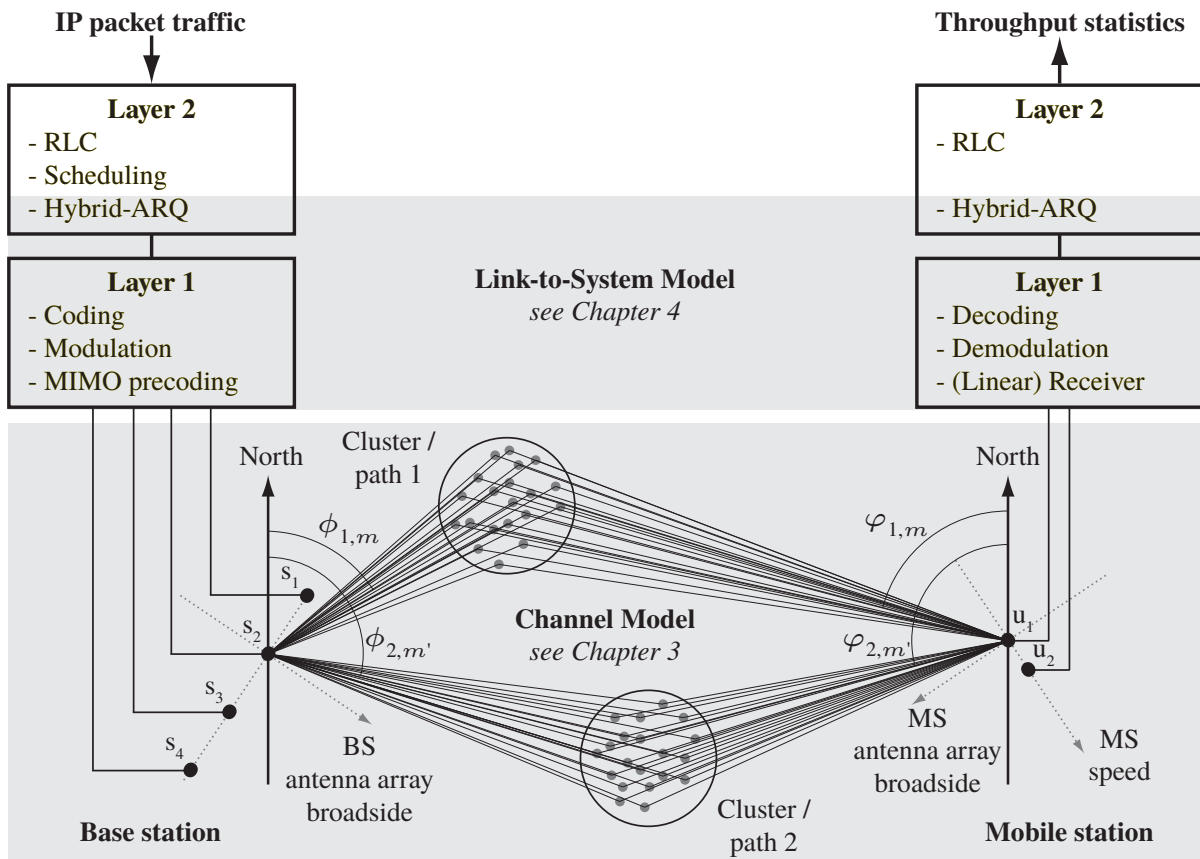


Figure 2.14.: Simplified view of downlink transmission through LTE protocol stacks and over MIMO channel showing the scope of system-level simulations

contrast to link-level simulations, which usually focus on a single communication link only, system-level simulations model all communication links in a system. Modeling all links allows to consider the mutual interference between multiple links, which is an important factor limiting the performance of a multi-cell system.

For system-level simulations, base stations and mobile stations are placed on a two-dimensional scenario that is generic, that is, not specific to a certain location. The users are dropped randomly into the scenario and for the duration of this *drop* they exchange packet-based data traffic in the downlink and uplink direction with their serving base station. Figure 2.14 shows IP packets traveling through a simplified protocol stack and over the wireless channel from one of the base stations to one of its associated mobile stations in the downlink direction. In a system-level simulation, multiple base stations at different sites each serve a number of mobile stations so that hundreds or thousands of such links are modeled. Depending on the underlying channel conditions of each link, and depending on how the radio technology handles them, a smaller or bigger number of packets correctly arrives at the receiving end of the link. The carried traffic (goodput) for each link, together with other metrics of interest like packet delays or SINRs, is included in the statistical evaluation of the system-level simulation.

The investigated radio interface technology (e.g., LTE) is modeled from the point where

the outgoing packets reach the layer 2 of the transmitter's protocol stack shown on the top of Fig. 2.4. IP packets are created by a traffic generator that models the behavior of higher layers. When determining spectral efficiencies, an inexhaustible traffic source keeps the senders' buffers filled at all times (*full buffer assumption*). Other evaluations, such as the determination of *voice over IP* (VoIP) capacities, make use of more sophisticated traffic models. For example, it is possible to distinguish between different QoS traffic classes and to model traffic sources with limited bit rates, whose packets might arrive with constant or variable bit rates. When discussing simulation results in the remainder of this thesis, we will always consider full-buffer traffic. For the same reason, we will also not derive any packet delay results. With full-buffer traffic, the IP packet delays would mostly depend on the assumed buffer sizes. In addition, apart from the short 1 ms subframe durations on the air, which can serve as a lower bound to the physical layer delays, most of the delay contributions are caused by the protocol stack, so that their duration depends on the implementation or on the modeling in a simulation.

Figure 2.14 shows a simplified protocol stack that is reduced to the layer 1 and 2 functionalities that are most relevant when evaluating the system spectral efficiency in system-level simulations. The figure also highlights the channel and link-to-system models, which we will discuss in detail in the following two chapters. The physical layer with its typical functions like channel coding, modulation, and MIMO signal processing is usually not modeled on the bit- or symbol-level in system-level simulations. This level of detail is only found in *link-level simulations* where, instead of a whole cellular system with hundreds or thousands of links, only a single link is considered. Instead, system-level simulations often employ a so-called *link-to-system model* (also called *PHY abstraction*) that aims at capturing how reliable the physical layer transports a data block from the sender to the receiver given the chosen resource allocation and the current channel and interference conditions (see Chapter 4 for details).

The channel model, together with the deployment scenario, forms the basis for link and system-level simulations. It models the general propagation conditions between base stations and mobile stations as well as their fluctuations in the time, frequency, and spatial dimension. Due to the temporal fluctuations of the channel and because traffic and thus interference conditions in each cell are potentially dynamic, system-level simulations are performed over a certain simulation time frame. During that time (typically, one or a couple of seconds), statistics are gathered to compute the metrics of interest. The channel conditions are not uniform over the scenario area so that each user's conditions depend on its randomly assigned position. As users do not move but are at fixed position during one drop, each user only samples the channel at its position. To achieve a statistically sound result that is representative for the whole scenario area, it is thus common to re-drop users to different positions during the course of the simulation. Usually, statistics are not gathered during the transient phase after a new drop to allow, for example, control loops to settle.

2.5. System-Level Simulation Tools

In the previous section, we motivated the use of simulations for assessing the performance of RRM schemes and the spectral efficiency of cellular systems. In this section, we will discuss several existing simulation tools and give a short overview of the IMTAphy IMT-Advanced channel model, link-to-system model, and LTE system-level simulator implementation that was developed by the author to carry out this work. By making the full source code of IMTAphy freely available as open-source on the Internet, we support the idea of *reproducible research* [VKV09]. System-level simulations are inherently complex so that it is not feasible to exhaustively describe all relevant algorithmic details and parameter settings under which results have been obtained. Providing the full source code thus allows all interested parties to reproduce the claimed results and inspect the implementation in detail.

2.5.1. A Short Survey of Simulation Tools

Before presenting our IMTAphy implementation, we will give a short overview of other simulation tools. Important criteria are timely availability, calibration against external references, focus on aspects relevant for the evaluation of system spectral efficiency, computational efficiency, as well as open-source availability. As we will see, none of the freely available open-source tools provides (all of) these properties which motivated the development of IMTAphy.

There are different classes of simulation tools. *Link-level* simulators are single-user and focus on physical layer topics like channel estimation, equalization, MIMO detection, channel coding, etc. Due to the detailed modeling, only a single serving BS-MS link and potentially a few interfering links can be simulated with reasonable computational effort. While the individual links can be modeled quite accurately, link-level simulations fail to capture realistic interference and multi-user scheduling constellations, especially if interdependencies between the resource allocation of different links exists. *System-level* simulators use an abstraction technique for modeling the transceiver performance (*link-to-system model*, see Chapter 4) and are thus able to simulate many more links and to model the interference between multiple cells and users. The focus of system-level simulation is on cell spectral efficiency and RRM schemes in the layers 1 and 2 of the protocol stack. Also in a multi-user and multi-cell setting but with an even higher degree of abstraction from the physical layer, *network-level* simulations focus on protocol stack signaling with the core network and between base stations as it is necessary, for example, for handovers.

System-level simulations, which are in the focus of this thesis, have to cope with a high degree of complexity. Not only multiple RRM schemes and their interactions have to be modeled, but the modeling has to be performed for hundreds to thousands of users in, for the typical example scenario, 57 cells. In addition, to capture the interference between all base stations and all users, tens of thousands of BS-MS links have to be considered. Due to these high computational demands, system-level simulators are

often implemented in compiled languages such as C++. Link-level simulations, on the other hand, oftentimes use MATLAB [MAT13] because they have a higher signal processing focus. Both system-level and link-level simulation tools need a radio channel model. Some simulators (e.g., IMT_{Aphy}) come with an integrated channel model implementation, while other tools make use of external channel models.

Channel Model Implementations

There is a wide variety of channel models used for simulations in the literature. Different models and their implementation vary with respect to their modeling accuracy and their computational complexity. As described in Section 2.1.1, the pathloss and shadowing are two major effects for which a large amount of empirical models exists. In the following, we will focus on the fast fading.

Due to the discrete nature of computer simulations, channels and signals are only generated for discrete points in time. If for multiple discrete signal realizations the radio channel realization is not updated, the channel models exhibit a so-called *block-fading* nature. Channel models can differ in the extension of these blocks, during which the fading is assumed to be constant, and in the way correlations between the channel realizations of adjacent bins in time, frequency, and space are generated. One of the simplest “channel models” for a SISO link with the system model introduced in (2.11) assumes a static channel transfer function (i.e., no fading, $h = \text{const}$) and an AWGN noise component n that is randomly generated for each realization following a circularly-symmetric complex Gaussian distribution. If also the channel coefficient h is generated this way with an expected power gain that is determined by the pathloss and shadowing attenuation, we have a Rayleigh-fading channel with uncorrelated fading in time and frequency.

In order to model correlations in the time and frequency domain, some models define a fixed *power delay profile* (PDP) for multiple paths that each exhibit independent Rayleigh fading generated according to a specified Doppler spread. Well-known *multipath channel models* are the *Vehicular A/B* and *Pedestrian A/B* models defined by the ITU in [M.1225] for the evaluation of IMT-2000 systems such as UMTS. These models, and also their LTE extensions *extended pedestrian A* (EPA), *extended vehicular A* (EVA), and *extended typical urban* (ETU) [36.101, Sec. B.2], define fixed relative expected powers and delays between the multipath taps c_i in (2.9). The individually fading path coefficients c_i are generated according to Jake’s model [Jak74; ZX03] with pre-defined Doppler frequencies of, for example, 5, 30, 70, or 300 Hz [36.101, Sec. B.2] that determine the temporal nature of the fast fading. If a MIMO channel matrix H is to be modeled, in addition the correlations between the antenna pairs play an important role as discussed in Section 2.1.2. A popular model to “artificially” introduce such correlations is the *Kronecker* model [KSP+02]. It establishes the desired degree of correlation by multiplying a channel matrix (consisting of independently generated fading realizations for each antenna pair) with the square roots of two (properly normalized) correlation matrices representing correlations introduced at the Rx and Tx side of the link. For example, [36.101, Sec. B.2.3] defines MIMO correlation matrices for “low”, “medium”,

and “high” correlation cases. A typical channel profile for LTE conformance testing could thus be called, for example, “EPA5 2 × 2 Low” defining a 2-by-2 MIMO channel with an extended pedestrian A PDP profile, a 5 Hz maximum Doppler frequency, and a low spatial correlation created with the Kronecker model.

The aforementioned channel modeling techniques are simple to implement in simulators and can be used for conformance testing and link-level simulations of isolated links. They are, however, not useful to characterize system-level performance, which depends on the characteristics of many links in a system. Besides pathloss and shadowing, also fast fading and MIMO propagation characteristics are correlated between the links from one user to multiple base stations and between users that are located in the same geographical area. In addition, the MIMO channel correlation depends on both the angular spread and the layout of the antenna array. The latter can be changed so that channel models that additionally take the angular characteristics of multipath propagation into account also allow to evaluate different antenna setups. A class of such channel models is called *geometry-based stochastic channel models* or *spatial channel models*. They include pathloss and shadowing models and generate the fast fading MIMO channels based on a *sum of sinusoids* approach where the fading coefficients for each path in the CIR (2.9) result from the superposition of rays reflected by reflection *clusters* as shown for an example with two such clusters in Fig. 2.14. The clusters are placed at random locations (“geometry-based”) defined by scenario-specific distributions of arrival and departure angles and also the PDPs are chosen randomly in a similar way (“stochastic”). Prominent examples of such channel models are the SCM [25.996] and SCME [BHS05] models, the WINNER models [KMH+08], and finally the IMT-Advanced model [M.2135], which is a special parametrization of the WINNER 2 model. MATLAB implementations of the IMT-Advanced channel model are available from the ITU’s website [Fin09b] but do not incorporate the [M.2135] deployment scenarios that will be discussed in Chapter 3. For the use in `IMTaphy`, we provide an efficient C++ implementation of the IMT-Advanced channel model, which we will present together with a detailed calibration campaign [Ell11] in Chapter 3.

In 2013, an evolved version of the WINNER channel models was published [JBTJ12] together with MATLAB source code that overcomes some limitations of the IMT-A models by supporting, for example, the movement of users with changing propagation conditions and a better antenna modeling. The 3GPP also studies [RP-122034] an extension of the current spatial channel models to support the evaluation of 3D beamforming extensions in future LTE releases.

Another option are *deterministic channel models*. Channel traces computed with *raytracing* methods or recordings captured from channel sounding campaigns in the field can offer the highest degree of realism. Their applicability, however, is limited to the specific site for which they were computed or measured and, in the case of raytracing, depends on the availability and quality of the building and terrain models that are used. Furthermore, due to the storage requirements, only a very limited amount of links or a short time duration is available to be used in simulations [Bon06].

Open-Source System-Level Simulation Tools

Most likely, all companies and organizations actively taking part in the standardization process have their own proprietary implementations of the channel model and the abstraction techniques presented in this thesis as well as a model of the LTE protocol stack. They use their simulators to contribute simulation results to standardization meetings, but they do not share simulation tools and there is no publicly available 3GPP system-level simulator reference implementation. In the following, we will review some LTE simulation tools that are available to the research community as open-source software. Commercial tools, of which there are a few, are not considered here because they are not suitable for academic research. The license costs are one reason, but the biggest drawback is that we cannot accept *black-box* implementations of important system components whose performance is to be studied.

The **TU Vienna** provides [MCŠ+11] a collection of link-level [MWI+09] and system-level simulators for LTE and LTE-Advanced. Their system-level simulator [IWR10; Iku13] uses a link-to-system mapping approach that is very similar to the one used in IMT_{Aphy}. Compared to IMT_{Aphy}, it does not come with an IMT-Advanced channel model implementation but allows to use the external WINNER MATLAB code. As mentioned above, the WINNER implementation [Fin09b] does not offer deployment scenarios according to the IMT-Advanced guidelines. The biggest limitation, however, is the choice of MATLAB which restricts the use of the Vienna simulators to smaller scenarios sizes with a limited number of links that can be modeled in detail.

The **openWNS** [BMKS09; MBS10] simulator is a modular simulation environment for wireless networks. It provides a simulation framework and a library for building protocol stacks [SDO+05]. Both are used by our IMT_{Aphy} implementation. Besides that, and completely independent of IMT_{Aphy}, there exists a separate LTE model that allows system-level simulations of LTE systems to a certain extent. The corresponding channel model is calibrated against the IMT-Advanced large-scale propagation models (pathloss and shadowing, see Chapter 3), but it does not offer a suitable fast-fading MIMO channel model. Hence, no MIMO RRM schemes can be evaluated.

Another implementation in C++ is the **LTE-Sim** simulator [PGB+11], which has more of a network focus as it supports mobility and handovers. It was used by its creators for the evaluation of various scheduling algorithms in a survey publication [CPG+13]. The presented results are, however, only of limited relevance because no MIMO transmission modes are supported and the channel modeling is based on a simple Jake's model for fast fading.

The **LTE-EPC Network Simulator (LENA)** [BMRN11] was first released in 2011 as a module for the ns-3 [NS3] network simulator. Even though the authors claim that it supports the evaluation of RRM schemes such as scheduling and interference coordination, it lacks realistic channel models and only models LTE MIMO transmission schemes by relying on pre-determined [CGE03] gains associated with different MIMO modes. This approach neglects performance impacts of correlated channels and of the spatial interference structure. Also, no calibrated deployment scenario setups complying with 3GPP [36.814] or ITU-R [M.2135] guidelines are provided that would allow to

derive comparable results. The strengths of LENA, however, are in the detailed modeling of the LTE protocol stack and the EPC integration. As the name suggests, it belongs more into the network simulator rather than the system-level simulator category.

2.5.2. Overview of IMTAphy

All of the system-level simulation tools available as open-source software have their weaknesses so that for this thesis, an independent tool was developed by the author. The IMTAphy simulation tool has three main components, all of which are implemented in C++: an IMT-Advanced channel model implementation, a link-to-system model, and a simplified model of the BS and MS LTE user-plane protocol stacks featuring various RRM schemes for the evaluation of, for example, MIMO scheduling and link adaptation algorithms in LTE. In early 2011 the IMT-Advanced channel model implementation was released as open-source software and its calibration was described in [Ell11]. The link-to-system model and first LTE calibration results were then presented by the author [EllPres11] in July of 2011. IMTAphy has been released on the Internet [IMTAphyLP; IMTAphyWeb] under the GNU Public License [Fre07]. Since the publication of the full simulation tool at the end of 2011, the IMTAphy web page [IMTAphyWeb] has attracted more than 5400 visitors until April of 2014.

To the best of the author's knowledge, until the submission of this thesis, IMTAphy remains the only freely available implementation of the IMT-Advanced channel model and deployment scenarios. One C++ [TTA] and one MATLAB [Fin09b] implementation of the channel model are available from the ITU's web page, but neither of them is fully calibrated against the ITU guidelines. In comparison to IMTAphy, both implementations are computationally inefficient.

IMTAphy is not a standalone simulator but requires the openWNS [BMKS09] simulation framework, which is also available as open-source. Written in C++ and featuring a configuration framework [BMKS09] based on *Python*, the published source code of the IMTAphy module encompasses in excess of 37 000 lines of code excluding comments. According to the independent website `ohloh.net`, which analyzes open-source software projects, the estimated programming effort, derived using the basic version of the COCOMO software cost estimation model [BCH+95], amounts to 9 person-years or half a million US dollars [Oh13].

Simulator Software Architecture

The IMTAphy simulator features an object-oriented design. The run-time components are written entirely in C++ and make extensive use of vector math libraries [MKL] to speed up the computation. In addition, some computations are performed in a multi-threaded way using the openMP [OMP] parallel programming library. The IMTAphy C++ classes have well-defined interfaces that allow to exchange different derived variants and to combine them flexibly. For example, different schedulers or link adaptation

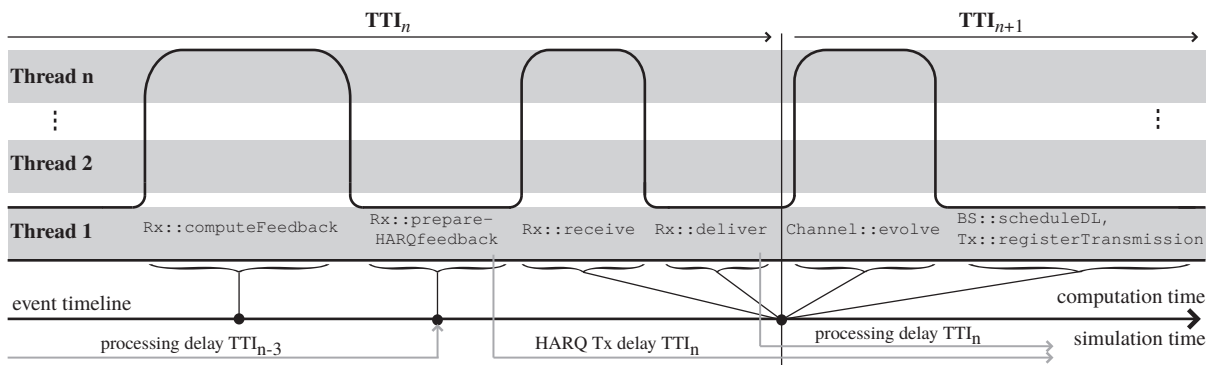


Figure 2.15.: Simulation flow in IMTAphy with multi-threading of selected functions

algorithms each share a common interface base class that allows to plug different variants of these RRM schemes into the protocol stacks without any change to the remaining system. The runtime instantiation is controlled by a configuration mechanism that relies on Python scripts that specify which (specialized) class is to be instantiated and how each configuration parameter is to be set. That way, a huge amount of configuration options (deployment and channel model properties, RRM algorithms, parameters, etc.) can be controlled at runtime in order to evaluate the performance of different system configurations. Each base station and each mobile station is modeled as an object that contains further objects forming the station's protocol stack. In addition, some central objects control simulator-wide mechanisms like, for example, the channel model, the simulation event processing, and the statistical evaluation at the end of the simulation.

The openWNS simulation framework [BMKS09] provides functionalities to setup simulation campaigns consisting of many individual simulation jobs, each with a specific combination of parameter settings. A compute cluster at the Institute for Communication Networks allowed performing up to 100 simulation jobs in parallel. With the help of the openWNS tools, the results of each run are exported to a central database, which can be accessed with a graphical user interface that allows to visualize simulation results by aggregating the results from many different individual simulation jobs. In total more than 200 000 simulation jobs were performed for this thesis.

Simulation Flow

Figure 2.15 illustrates the simulation flow in IMTAphy. The general concept followed by the openWNS simulator on which IMTAphy is based, is that of a *time-discrete event-based stochastic simulation*. *Time-discrete* means that the *simulation time* jumps in discrete steps between points in time when some *event* has to be processed, for example, because a new packet is generated or an LTE transmission time interval starts. The simulation is *stochastic* because many actions depend on random numbers: the inter-arrival time of packets, the initialization of small-scale parameters like delays or angles in the channel model (see also Chapter 3), the successful reception of a transport block that has some associated block error probability, etc. The required random numbers

in `IMTAphy` are generated using a single *Mersenne Twister* [MN98] random number generator provided by the *boost* C++ libraries [boost]. The initial seed for the random number generator can be configured at runtime so that multiple simulation runs with identical parameters, but independent random numbers, can be performed and averaged to obtain more representative results.

Apart from randomly generated packet arrival times, most events in the `IMTAphy` simulation flow follow the LTE time structure as shown, for example, in Fig. 2.7. Figure 2.15 illustrates the corresponding execution flow in `IMTAphy` between some TTI n and the following TTI $n + 1$. During the TTI (and depending on the feedback frequency), the MS receivers compute the CQI, PMI, and RI feedback according to the current channel and interference conditions. Since this computation can happen independently in each of the (potentially hundreds of) MS receivers, it is parallelized on the MS receiver level. At the border of two TTIs, most of the processing happens. As Fig. 2.15 exhibits, a number of consecutively processed events correspond to this single point in time. First, the transmissions during the just ended TTI n are processed by the link-to-system model described in Chapter 4. This can happen in parallel because no interdependencies between receivers exist. Afterwards, the results of this step are delivered to the PHY interface in the receivers. This occurs sequentially because in order to model the delays arising from the Turbo decoder processing, each receiver creates an event for the point in time when the processing is assumed to be over. The event scheduler in `openWNS` is not thread-safe so that the scheduling of this event has to be done in a single-threaded manner. The example in Fig. 2.15 shows that during TTI n a previous processing delay expired which triggers the preparation of HARQ feedback based on the final ACK/NACK verdict of the received transport block. For the HARQ feedback, another timer is programmed in order to model the delay with which the HARQ feedback is made available at the transmitter. Like all other feedbacks (i.e., CQI, PMI, and RI), the feedback transmissions over the radio channel are assumed to be error-free and are not modeled in detail except for the time delay and quantization according to the standard. In the next event at the beginning of TTI $n + 1$, the channel model is triggered to evolve the radio channels to the next TTI's time. As described in Chapter 3, the time evolution of the channel model is a highly optimized process that exploits vectorization and multi-threading features of modern CPUs. Again, the channel model time update can happen in parallel for each of the links. Finally, the scheduling process in each base station is triggered. It is performed in a single-threaded way to avoid *race conditions* if two scheduler threads try to simultaneously access a common variable (e.g., for BS coordination schemes).

Depending on the size of the scenario, the number of modeled MIMO antenna pairs, and the complexity of the investigated algorithms, the computing time per unit of simulation time varies. Being a system-level simulator that models dozens of base stations and hundreds to thousands of users at once, `IMTAphy` is not designed to perform in *real time*. Most simulation setups need less than one hour of computation time for one second of simulation time so that the required simulation times of a few seconds can be achieved in a few hours even on the hardware from 2008 available in the compute cluster of the institute.

2.6. Summary and Contributions

In this chapter, we introduced key terms and concepts used in the remainder of the thesis. We discussed how analyzing the downlink system spectral efficiency of a radio access technology like LTE differs from more information-theoretic approaches. Our focus is on the data rates that can be achieved in an actual real-world LTE multi-cell system operating according to the LTE standard.

To this end, we gave a brief introduction to the main aspects of the LTE protocol stack and physical layer that determine the link spectral efficiency, that is, the possible rates considering an isolated BS–MS link. The LTE transmission modes determine how the link can be operated, for example, what kind of MIMO processing and how much channel state feedback is available. The choice of modulation and coding schemes directly limits the data rates on a link. The control channel and reference signal overhead is another aspect we discussed in detail. Besides these features of the standard that influence the link spectral efficiency, one important aspect of the link spectral efficiency is the receiver architecture. We briefly discussed the differences between MLD-like schemes and the linear receiver, which we will use throughout the thesis due to its importance and well-understood modeling.

What distinguishes a per-link analysis from the system-level approach followed in this thesis, is the multi-user and multi-cell aspect. On the one hand, the system spectral efficiency reflects the statistical distribution of per-user radio conditions like pathloss and shadowing throughout the cells of the system. On the other hand, channel-adaptive schedulers aim at increasing the system spectral efficiency by exploiting the multi-user diversity of the fast fading between the users in the cell. On top of that, all cells are linked by inter-cell interference with the strength and the structure of the inter-cell interference intrinsically limiting the data rates in a cellular systems.

Due to this complexity and the intention to evaluate the performance of concrete RRM schemes like feedback signaling, scheduling, spatial processing, and link-adaptation strategies, system spectral efficiency evaluations are typically performed by computer simulations. To enable comparable system-level simulations, the ITU proposed simulation guidelines within its IMT-Advanced framework. Motivated by the lack of suitable freely-available simulation tools, we developed an LTE system-level simulator called `IMTphy`. It contains an efficient and fully-calibrated implementation of the IMT-Advanced channel model (see Chapter 3), a calibrated state-of-the-art link-to-system model (see Chapter 4), and is available as open-source on the Internet enabling the reproduction of all system spectral efficiency results presented in this thesis.

The eponymous contribution in this chapter is the analysis of the so-called PHY losses that separate the actual achievable spectral efficiency in an LTE system from the channel capacity. Further, a review of system-level simulation tools and channel model implementations was presented. Another contribution is the discussion of an efficient multi-threaded software architecture for LTE system-level simulation. Its implementation in `IMTphy` allowed performing more than 200 000 simulation runs to obtain the results for this thesis.

3. The IMT-Advanced Channel Models

In this chapter, we introduce the IMT-Advanced deployment scenarios and channel models, which are the current standard framework for performance evaluation of cellular systems. We do not contribute own models here but adopt all IMT-Advanced assumptions in order to provide a basis for comparable performance evaluation. We do provide an efficient open-source implementation of the IMT-Advanced models in our `IMTAphy` simulator and also discuss the detailed calibration of our implementation against results from industry-leading companies and organizations. A part of the material presented in this chapter has been published by the author in [Ell12a].

3.1. Test Environments and Deployment Scenarios

Figure 3.1 provides an overview of the deployment parameters on the BS and MS side that are in the focus of the first part of this chapter. The ITU Radiocommunication Sector (ITU-R) report [M.2135] distinguishes between so-called *test environments*, the corresponding *deployment scenarios*, and the associated *channel models*. Table 3.1 lists the deployment scenarios that have been selected by the ITU for the IMT-Advanced evaluation: the *indoor hotspot scenario*, the *urban micro-cell scenario*, the *urban macro-cell scenario*, and the *rural macro-cell scenario*. They belong to the *indoor*, the *microcellular*, the *base coverage urban*, and the *high speed* test environment, respectively. For each deployment scenario there is a channel model of the same name. In this text, we will refer to the different deployment scenarios and channel models by using the acronyms (InH, UMi, UMa, and RMa) for the channel models as introduced by the ITU [M.2135]. The aforementioned scenarios are considered mandatory by the ITU. In addition, there is an optional *suburban macro* (SMa) scenario and channel model that belongs to the base coverage urban test environment. It is fully specified in report [M.2135], but we will not discuss it here.

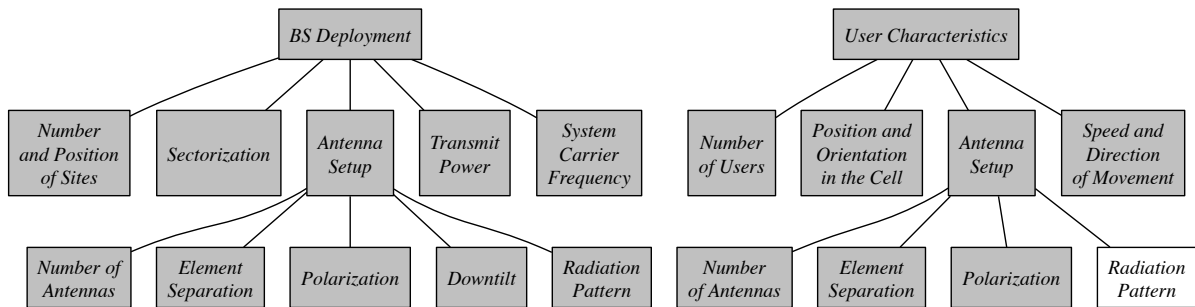


Figure 3.1.: Overview of the main BS and user characteristics covered in this chapter

3. The IMT-Advanced Channel Models

Table 3.1.: Overview of deployment scenarios and simulation parameters [M.2135]

Parameter / deployment scenario	Indoor hotspot	Urban micro-cell	Urban macro-cell	Rural macro-cell
Test environment	indoor	microcellular	base coverage urban	high speed
Channel model	InH	UMi	UMa	RMa
Total BS Tx power (20 MHz)	21 dBm	44 dBm	49 dBm	49 dBm
Max. MS transmit power P_{Max}	21 dBm	24 dBm	24 dBm	24 dBm
Center frequency f_c	3.4 GHz	2.5 GHz	2 GHz	800 MHz
BS height h_{BS}	6 m	10 m	25 m	35 m
BS antenna downtilt θ_{tilt} [36.814]	0°	-12°	-12°	-6°
Number of sites	2	19	19	19
Sectorization	none	3 sectors	3 sectors	3 sectors
Inter-site distance (ISD)	60 m	200 m	500 m	1732 m
Min. BS-MS distance	3 m	10 m	25 m	35 m
User locations	indoor	50% indoor, 50% outdoor	in vehicle	in vehicle
User speed v	3 km/h	3 km/h	30 km/h	120 km/h
Coherence time based on (2.13)	26.5 ms	36.0 ms	4.5 ms	2.8 ms

Most of the deployment scenario assumptions presented in Tables 3.1 and 3.2 and in the following sections are based on the original ITU guidelines [M.2135] and subsequent “Guidelines for using IMT-Advanced channel models” [Fin09a]. These were issued to help the evaluation groups interpret the original guidelines. Some other parameters such as the BS antenna downtilt have not been explicitly specified by the ITU. For these we will use the 3GPP’s assumptions for system-level simulator calibration as described in the annex of technical report (TR) [36.814].

All deployment scenarios except the InH scenario use a conventional cellular deployment model with 19 base station sites. Three base stations, each equipped with a sectorized antenna array serving a hexagonal cell, are co-located per site as depicted in Fig. 3.2. The different hexagonal scenario types all consist of $19 \cdot 3 = 57$ cells in total.

Table 3.2.: Common deployment scenario parameters [M.2135]

Parameter	Value
BS noise figure	5 dB
MS noise figure	7 dB
BS antenna gain	17 dBi (0 dBi for InH)
MS antenna gain	0 dBi
Avg. users per cell	10
Feeder loss at BS	2 dB for large-scale calibrations, 0 dB otherwise [36.814]
Handover margin	1 dB [36.814]




Figure 3.2.: Wraparound configuration surrounding a 21 cells (7 sites) scenario with 6 shifted copies of itself; arrows mark the offset of each copy

But the cell sizes and thus the covered area vary greatly due to different inter-site distances (ISDs) that range from 200 m (for UMi) to 1732 m (for RMa), see also Table 3.1. The InH scenario, in contrast, is a very small scenario with just two base stations equipped with omnidirectional, that is, non-sectorized antennas. The base stations are located at 60 m ISD in the middle of a small $120 \text{ m} \times 50 \text{ m}$ indoor area. Table 3.1 lists further differences between the deployments like different BS and MS transmit power budgets, which refer to 20 MHz of spectrum. Several parameters are common to all deployments, see Table 3.2. One main difference in deployment characteristics is the assumed carrier frequency and user mobility velocity. Together, they determine the channel coherence time. Only the InH and UMi scenarios have long-enough coherence times to fully benefit from channel-adaptive scheduling, as we will see in Chapter 5.

For each simulation drop, $10 \cdot 57 = 570$ users are randomly and uniformly distributed over the whole area of the hexagonal scenarios and $10 \cdot 2 = 20$ users are dropped onto the rectangular indoor hotspot area. The user positions are kept fixed during the drop duration. In order to model the fast-fading channel, the users are imagined to move at a scenario-specific speed into randomly assigned directions. A BS serves 10 mobiles on average. Each MS is always only associated to the base station of one cell that is called its *serving cell*. The exact amount of users per BS depends on the random spatial distribution of the users, on the random shadowing (see below), and on the association and handover process of the radio interface technology. The handover process itself usually cannot be simulated because the MS positions are fixed and only the fast-fading varies during the drop. But to account for the handover hysteresis, which keeps a mobile associated to its previous BS while another is already stronger, the 3GPP assumes [36.814] that the mobiles are randomly associated to one of the base stations whose *reference signal received power* (RSRP) level is within a 1 dB *handover margin* from that of the BS with the strongest received signal.

Many of the channel model properties presented in Section 3.3 depend on distances and angles between base station sites and mobiles. All distances and angles are measured in the xy -plane in two dimensions only without taking BS or MS heights into account. The only exception are the antenna field patterns, which are modeled in three dimensions (see Section 3.2.1), so that the elevation angles between mobiles and base stations are computed from their relative heights and distances. To assure a minimum pathloss and coverage by the BS antenna elevation pattern, users that are too close to a base station site are re-dropped to a different position in the system area (see the min. BS–MS distance in Table 3.1).

Real-world cellular systems cover large areas with thousands of potentially interfering cells. They are typically inter-cell interference-limited. For modeling cellular systems, a good trade-off between a low complexity and an accurate interference level is to surround the three cells of a site under consideration by two rings of interfering sites ($6 + 12 = 18$ sites with $3 \cdot 18 = 54$ interfering cells). The interference from cells farther away can be neglected due to the high pathloss. However, from the 57 cells in total, only the performance of the three inner cells could be evaluated because the remaining cells are more or less on the edge of the scenario area where they do not receive equally strong interference from all directions. To avoid this, a so-called *wraparound* technique can be used. It creates a torus-like deployment topology where all cells appear to be surrounded by two rings of interfering sites. In this way, all cells face a realistic and uniform inter-cell interference situation and can all be included in the performance evaluation. For IMT-Advanced system-level simulations wraparound modeling is mandatory [M.2135]. In *IMT_Aphy*, we realize this on a per BS–MS link basis by shifting an MS to that position among 7 possible positions that yields the smallest Euclidean distance in the xy -plane to the BS. In other words, for each station we model a link to come from the direction of the shortest distance and not from all possible directions at once. Figure 3.2 shows an exemplary 21 cells scenario (highlighted cells in the center) surrounded by 6 virtual copies, which are shifted from the scenario’s center according to the plotted shift vectors. Of course, a wraparound is not applied in the InH scenario because it models two isolated base stations in an indoor setting.

3.2. Antenna Modeling

The IMT-Advanced channel model introduced in the following section separates the modeling of BS and MS antennas from the modeling of the spatial propagation channel itself because there is a fundamental difference: the propagation channel, on the one hand, is caused by the environment and can only be accurately modeled, but it cannot be changed. The antennas, on the other hand, can be designed by a system engineer. Thus, it is important to support arbitrary far-field antenna patterns and antenna array layouts in a system-level simulator.

In Section 3.2.1, we will first present how the field pattern of each antenna element is modeled within the IMT-Advanced guidelines. Then, in Section 3.2.2, we will in-

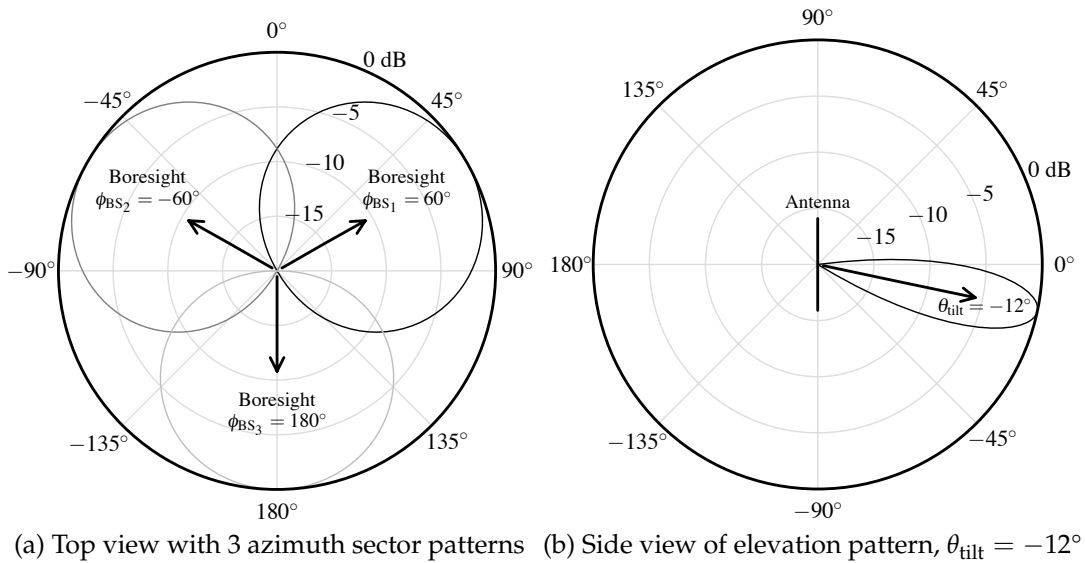


Figure 3.3.: Polar plots of antenna gain patterns

roduce the antenna array layouts (i.e., the separation between multiple antenna elements and their slants) used for IMT-Advanced simulator calibration and general performance evaluation of 3GPP LTE and LTE-Advanced.

3.2.1. IMT-Advanced Antenna Field Pattern Model

The MS antennas, as well as the InH BS antennas, have an omnidirectional pattern with 0 dBi gain. The BS antennas for the hexagonal scenarios should only radiate into a 120° wide sector and also focus the radiated power downwards into their own cell. The ITU thus specifies a 3D pattern whose azimuth (in the xy -plane) and elevation (upward from the xy -plane) power gain patterns are shown as polar plots in Fig. 3.3.

The 3D antenna pattern is defined [M.2135] as having a maximum gain $G(\phi, \theta) = 0$ dB (in addition to a 17 dBi gain in the boresight direction) and a forward-backward attenuation of $G_{\text{Min}} = 20$ dB resulting in a maximum attenuation of $G(\phi, \theta) = -20$ dB. The pattern is defined in two components with respect to the azimuth angle ϕ as $G_{\text{Azimuth}}(\phi)$ in (3.1) and with respect to the elevation angle θ as $G_{\text{Elevation}}(\theta)$ in (3.2). As can be seen in Fig. 3.3, the elevation pattern with a 3 dB beamwidth $\theta_{3\text{dB}} = 15^\circ$ is much narrower than the azimuth pattern with a 3 dB beamwidth $\phi_{3\text{dB}} = 70^\circ$. The final 3D pattern is computed as the sum of the attenuations $G_{\text{Azimuth}}(\phi)$ and $G_{\text{Elevation}}(\theta)$ by applying (3.3) in a way that assures that the maximum attenuation of $G_{\text{Min}} = 20$ dB is not exceeded. The patterns are defined relative to the BS antenna array broadside direction azimuth angle ϕ_{BS} and the downtilt angle θ_{tilt} . The downtilt is not specified by the ITU so that Table 3.1 lists the 3GPP default values.

$$G_{\text{Azimuth}}(\phi) = -\min \left[12 \left(\frac{\phi - \phi_{\text{BS}}}{\phi_{3\text{dB}}} \right)^2, G_{\text{Min}} \right] \text{ in dB, } -180^\circ \leq \phi - \phi_{\text{BS}} \leq 180^\circ \quad (3.1)$$




Figure 3.4.: Antenna setup for calibration simulations defined in TR [36.814]

$$G_{\text{Elevation}}(\theta) = -\min \left[12 \left(\frac{\theta - \theta_{\text{tilt}}}{\theta_{3\text{dB}}} \right)^2, G_{\text{Min}} \right] \text{ in dB}, \quad -90^\circ \leq \theta - \theta_{\text{tilt}} \leq 90^\circ \quad (3.2)$$

$$G(\phi, \theta) = -\min \left[- (G_{\text{Azimuth}}(\phi) + G_{\text{Elevation}}(\theta)), G_{\text{Min}} \right] \text{ in dB} \quad (3.3)$$

3.2.2. 3GPP Antenna Array Layout and Polarization Model

MIMO communication systems inherently need multiple antennas at both ends of the wireless link. Together with the channel itself, the properties of the antenna array layouts at both sides of the link determine the correlation properties between the transmitter's output and the receiver's input signals as described in Section 2.1.1. Roughly speaking, MIMO channels with low correlation allow single-user spatial multiplexing transmissions, while channels with high correlation are suitable for beamforming or MU-MIMO transmission modes. Consequentially, the antenna array layouts at both sides of the link have an effect on the feasible transmission modes in the system and should be chosen carefully.

Two main design parameters of the antenna arrays influence the correlation properties [Win98]: the spatial separation of the antenna element columns and the polarization of the individual elements. Whether a certain element spacing results in correlated or uncorrelated channel matrices depends on the angular spread of the incoming and outgoing multipath components of the signal. For a mobile station, which is usually surrounded by scatterers and thus experiences a high angular spread (see Fig. 3.11 for exemplary angular spreads), antenna spacings on the order of half the wavelength already contribute to uncorrelated channels. For base stations, which usually experience low angular spreads, the same $\lambda/2$ antenna element spacing results in high correlations. BS antenna elements for spatial multiplexing operation are thus usually separated by several wavelengths [Win98]. The multipath propagation and the angular spreads at the BS and MS will be discussed in more detail in Section 3.3.

Deploying antenna arrays with individual elements that are slanted so that they excite two orthogonal polarizations (i.e., vertical and horizontal) of the electrical field, is another way to contribute to low correlation channel matrices [Win98]. As an additional advantage, two slanted elements can be co-located so that the array can be built in a compact way. However, Winters [Win98] notes that for elevated base station antenna positions, the horizontal polarization component can be 6–10 dB weaker than the vertical polarization.

Figure 3.4 shows the BS and MS antenna array layout that is used by the 3GPP for their IMT-Advanced simulator calibration campaign. At both ends, two antennas are




Figure 3.5.: BS antenna configurations A) to E) with four antenna elements for IMT-Advanced evaluation simulations as defined in TR [36.814]

assumed with purely vertical polarization. The 10λ antenna spacing at the BS contributes to a low correlation channel matrix. In Fig. 3.5, the different BS antenna array configuration options A) to E) defined by the 3GPP in TR [36.814] for LTE-Advanced performance evaluation are shown. Configuration A), with a large element spacing, is suitable for spatial multiplexing, whereas configuration C) is fitting for rank-1 precoding or MU-MIMO operation. Configuration B) could be used for transmission modes that serve two users with two layers each because the closely-spaced group should allow to spatially separate users while the large spacing between the groups should contribute to uncorrelated channels for two-layer spatial multiplexing to each user. Configurations D) and E) employ slanted antenna elements, which have electrical field pattern components in both the horizontal and vertical polarization direction. The independent fading between polarization directions makes also these setups attractive for spatial multiplexing operation.

The spatial channel model introduced in Section 3.3 distinguishes between horizontal and vertical polarization components of the antenna field pattern. However, the IMT-Advanced evaluation guidelines [M.2135] only specify simple antenna patterns with purely vertical polarization. The 3GPP describes the modeling of polarized antennas with arbitrary polarization slant angles in the annex of TR [36.814] but only uses purely vertically polarized antennas for their calibration setups.

In TR [36.814] the 3GPP defines a simple polarized antenna model that models the horizontal and vertical components of the electrical field by a simple vector decomposition of the electrical field vector $F(\phi, \theta, \xi)$ into the horizontal $F^{(H)}(\phi, \theta, \xi)$ and vertical $F^{(V)}(\phi, \theta, \xi)$ components according to the slant angle ξ of the antenna element:

$$F^{(V)}(\phi, \theta, \xi) = \|F^{(V)}(\phi, \theta, \xi)\| = \sqrt{G(\phi, \theta)} \sin(\xi), \quad (3.4)$$

$$F^{(H)}(\phi, \theta, \xi) = \|F^{(H)}(\phi, \theta, \xi)\| = \sqrt{G(\phi, \theta)} \cos(\xi). \quad (3.5)$$

The electrical field strength is obtained directly from the antenna radiation power pattern (3.3) without considering that the horizontal and vertical components, as perceived from certain azimuth and elevation angles, might depend on these angles.

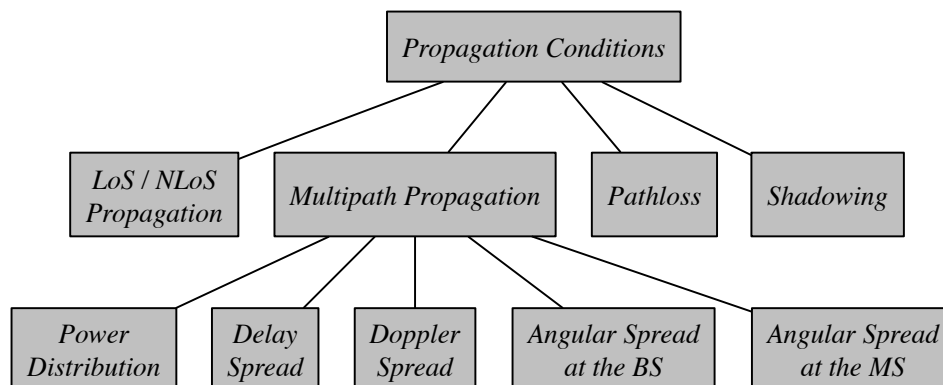


Figure 3.6.: Overview of propagation aspects covered by the IMT-A channel models

3.3. The IMT-Advanced Channel Models

After considering the assumptions on the BS and MS deployments, we now turn our attention to the IMT-Advanced channel models. Figure 3.6 gives an overview of the aspects that have to be modeled here. As introduced in Section 2.1.1, the main propagation effects are pathloss, shadowing, and the multipath propagation that leads to fast fading. Of course, the deployment scenarios and BS and MS characteristics, which we have discussed so far, influence the channel models and determine which parameters are used for the channel modeling.

For the evaluation of IMT-Advanced systems, the ITU defined [M.2135] a channel model consisting of a *primary module* and an *extension module*. The primary module consists of pathloss models, scenario-specific parameter sets for the InH, UMa, UMi, RMa, and SMa scenarios, and a so-called *generic model*. The generic model specifies the mathematical model and the algorithms used for channel modeling that apply to all scenarios. It is based on the WINNER 2 channel model [KMH+08] which, in turn, is similar to the 3GPP / 3GPP2 SCM channel model [25.996]. The scenario-specific parameters are used to instantiate scenario-specific spatial channel models. The optional extension module allows to generate alternate parameter sets for the UMa, RMa, and SMa scenarios. Here, we will focus on the primary module with its default parameter sets because this was the one used for the IMT-Advanced evaluations conducted by the 3GPP and the independent evaluation groups.

The channel model, as shown in Fig. 2.14 on page 49, models the propagation channel between each BS antenna s and each MS antenna u for a link between a base station and a mobile station. In a system-level simulation, the channel model has to cover all links, both serving and interfering. In fact, being able to model the different (but potentially correlated) channel conditions on different links and between different antenna pairs of the same link, is what makes the channel model suitable for MIMO system-level simulations.

The generic model employs a geometry-based stochastic approach to represent the multipath propagation channel between base stations and mobile stations. As an end

result, the model provides the time-variant *channel impulse response* (CIR) $c_{u,s}(t)$ for an antenna pair (u, s) according to (2.9) on page 11. For each of the N_{MPC} paths, a complex path coefficient $c_{u,s,n}(t)$ describes the change in amplitude and phase a complex baseband signal perceives when propagating along the considered path n that causes a delay of τ_n relative to the fastest path. Besides these multipath effects, the signal also suffers from large-scale fading effects such as pathloss and shadowing that directly scale the norm of the path coefficients $c_{u,s,n}$. Note that path coefficients, delays, and all related values generally differ between different BS–MS links. By applying a DFT as in (2.10) on page 11, the CIR can be transformed into the frequency domain resulting in the *channel transfer function* (CTF) $h_{u,s}(t, f)$ at a given frequency f .

The ITU generic model is not site-specific and as such does not model actual objects in the environment of a station like buildings that serve as scatterers and reflectors for the multipath components. Instead, it generates a number N_{MPC} of scattering clusters. Each cluster is assumed to reflect the signal on its way from the transmitter to the receiver and contributes one path to the multipath propagation. Two exemplary reflection clusters are shown in Fig. 2.14 on page 49. The geometric location of the n -th cluster is implicitly defined by its propagation delay τ_n and the angles under which cluster n is seen from the BS and MS. The nomenclature follows a downlink perspective so that the angles ϕ_n at the BS are called *angles of departure* (AoDs), while the angles φ_n at the MS are called *angles of arrival* (AoAs). The model is stochastic because the actual angles and delays are generated randomly from a properly configured probability distribution. With the AoAs, the AoDs, and the delays, the model captures the properties of multipath propagation that determine the fading nature of the MIMO channel as introduced in Section 2.1.1: In the frequency domain, the individual delays τ_n of the paths cause the *frequency-selective* fading. Wavefronts of multipath components arriving with distinct AoAs at the receive array are seen with different relative phase offsets at distinct receive antenna array elements $u_i \neq u_j$. This leads to a *spatially-selective* fading behavior.

Figure 2.14 on page 49 illustrates two exemplary paths between a base station and a mobile station antenna array. Depending on the chosen deployment scenario and on the propagation condition, between 10 and 20 clusters per link have to be modeled in the IMT-Advanced channel model. In addition, the clusters are assumed to have a substructure of $M = 20$ individual scatterers, as the figure shows. Each of these scatterers reflects a ray m under slightly different angles of departure $\phi_{n,m}$ and arrival $\varphi_{n,m}$ as seen from the BS and MS, respectively. Introducing such a per-cluster angular spread increases the angular resolution of the channel model which is important for accurate MIMO evaluations. In (3.17) on page 73 we will show how a single path coefficient $c_{u,s,n}(t)$ is computed from the superposition of these 20 rays per cluster. Before we get there, however, we will first introduce how the necessary large-scale and small-scale parameters are generated.

Table 3.3.: Line-of-sight propagation probability model and simulation results

Scenario	Probability as a function of distance [M.2135]	Per link probab.	
		All	Serving
InH	$\Pr(d) = \begin{cases} 1 & d \leq 18 \\ \exp((18-d)/27) & 18 < d < 37 \\ 0.5 & d \geq 37 \end{cases}$	67 %	92 %
RMa	$\Pr(d) = \begin{cases} 1 & d \leq 10 \\ \exp((10-d)/1000) & d > 10 \end{cases}$	12 %	82 %
UMa	$\Pr(d) = \min\left(\frac{18}{d}, 1\right) \left[1 - \exp\left(\frac{-d}{63}\right)\right] + \exp\left(\frac{-d}{63}\right)$	4 %	37 %
UMi	$\Pr(d_{\text{out}}) = \min\left(\frac{18}{d_{\text{out}}}, 1\right) \left[1 - \exp\left(\frac{-d_{\text{out}}}{36}\right)\right] + \exp\left(\frac{-d_{\text{out}}}{36}\right)$	9 %	75 %

3.3.1. Large-Scale Link Properties

As a first step, a set of so-called [M.2135] *large-scale parameters* is determined for each link in the system. Two broad classes of large-scale parameters exist. The first class, including pathloss and shadowing, describes the large-scale fading a link experiences. The other class, including the delay spread and angular spread of a link, serves to parameterize the probability distributions from which the individual delays and angles are drawn (the *small-scale parameters*).

Determining Line-of-Sight Propagation Conditions

After the users have been dropped into the selected deployment scenario, the first step to initialize the channel model is to randomly classify each BS–MS link to exhibit LoS or NLoS propagation. As Table 3.3 shows, depending on the two-dimensional (2D) distance d , a LoS probability $\Pr(d)$ is applied, reflecting that LoS propagation is more likely the closer a user is to a base station site. Note that the propagation condition is decided on a per-site basis so that a user always experiences identical propagation conditions to all base stations co-located at one site. Table 3.3 also lists the percentage of LoS propagation observed in system-level simulations. Except for InH, NLoS propagation dominates if all links (serving + interfering) are considered. Many serving links are LoS though because users are close to their serving base station, and having an LoS to a certain BS increases the chances that this BS is seen as the strongest BS and thus gets selected as the serving BS.

Depending on the propagation condition, different parameters are used for the pathloss and spatial channel model, as explained in the following sections. The UMi scenario is a special case because in UMi half of the users are outdoors and half of them are indoors. This classification is done randomly with a 50 % probability when a UMi drop is initialized. For the outdoor users, the propagation can be LoS or NLoS and the respective pathloss models and channel model parameter sets are used accordingly. For the indoor users, there is a special UMi *outdoor-to-indoor* (O2I) channel model parame-

ter set. The UMi O2I propagation does not exhibit LoS Rician fading, but both NLoS and LoS pathloss models can be applied to the outdoor portion d_{out} of the distance between the MS and the BS [Fin09a]. The 75 % LoS probability in Table 3.3 thus only refers to the pathloss models, while only 75 % of the outdoor users, that is 37.5 % overall, can perceive LoS Rician fading. The outdoor distance $d_{\text{out}} = d - d_{\text{in}}$ is computed by subtracting a uniformly randomly chosen indoor distance $d_{\text{in}} \sim \mathcal{U}(0, \min[d, 25])$ from the total distance d in meters for each O2I link.

Pathloss Models

The user locations are also relevant when computing the pathloss on a link, because for UMi indoor users a wall penetration loss $\text{PL}_{\text{tw}} = 20$ dB and an additional indoor pathloss $\text{PL}_{\text{in}}(d_{\text{in}}) = d_{\text{in}} \cdot 0.5$ dB/m are added to the basic outdoor pathloss PL_{basic} :

$$\text{PL}^{(\text{UMi}, \text{O2I})}(d) = \text{PL}_{\text{basic}}(d_{\text{out}}) + \text{PL}_{\text{tw}} + \text{PL}_{\text{in}}(d_{\text{in}}) \text{ in dB.} \quad (3.6)$$

In the UMa and RMa scenarios, all users are in vehicles (see Table 3.1) so that the basic outdoor pathloss $\text{PL}_{\text{basic}}(d)$ is adjusted by a log-normal distributed vehicle penetration loss PL_{Veh} with a mean of 9 dB and a standard deviation of 5 dB:

$$\text{PL}(d) = \text{PL}_{\text{basic}}(d) + \text{PL}_{\text{Veh}} \text{ in dB} \quad \text{with} \quad \text{PL}_{\text{Veh}} \sim \mathcal{N}(9, 5^2) \text{ in dB.} \quad (3.7)$$

Note that the vehicle penetration loss PL_{Veh} is specific to each vehicle so that for one user the same vehicle penetration loss is assumed, regardless of the destination BS. The presented pathloss adjustments were already mentioned in the initial 2008 ITU guidelines [M.2135s] but were subsequently clarified in 2009 [Fin09a]. In the InH scenario no adjustments to the basic pathloss are necessary because no outdoor walls are penetrated and the loss due to indoor walls is already considered in the basic pathloss formulas. Some simulations account for a feeder loss (in the cable connecting the BS antenna) of, for example, 2 dB which can simply be added to the overall pathloss.

The ITU adopted the basic pathloss $\text{PL}_{\text{basic}}(d)$ expressions from the WINNER channel model, which is based on measurement campaigns and literature results [KMH+08]. There are special pathloss formulas for each deployment scenario (InH, UMa, UMi, RMa, and SMa) and for both LoS and NLoS propagation types. For LoS propagation in the hexagonal scenarios, a dual slope pathloss model is used that assumes a higher pathloss exponent for distances above a certain break point distance d_{BP} . The pathloss models are valid for system center frequencies in the range of 2–6 GHz (450 MHz to 6 GHz for RMa) and allow to specify certain parameters like BS and MS antenna heights (all scenarios) as well as street widths and average building heights (UMa, RMa, and SMa). Due to this configurability, the pathloss formulas given by the ITU [M.2135] are lengthy. For that reason, we choose to only present the UMa pathloss model in a simplified form assuming standard values as given in Table 3.1 and in report [M.2135]. The basic pathloss $\text{PL}_{\text{basic}}(d)$ for a given distance d in meters and NLoS propagation is computed according to (3.8) with a pathloss exponent of 3.9. For LoS propagation (3.9) distinguishes between two different pathloss exponents

Table 3.4.: UMA large-scale parameter distribution parameters [M.2135]

Large-scale parameter	Mean μ		Std. dev. σ		Correlation distance d_{corr}	
	LoS	NLoS	LoS	NLoS	LoS	NLoS
Delay spread (DS) in $\log_{10}(\text{s})$	-7.03	-6.44	0.66	0.39	30 m	40 m
Angular spread in $\log_{10}(\text{degree})$ of:						
– departure angles (ASD)	1.15	1.41	0.28	0.28	18 m	50 m
– arrival angles (ASA)	1.81	1.87	0.20	0.11	15 m	50 m
Shadow fading (SF)	0 dB	0 dB	4 dB	6 dB	37 m	50 m
Rician K factor (K)	9 dB	n/a	3.5 dB	n/a	12 m	n/a

based on a break point distance threshold of $d_{\text{BP}} = 4(h_{\text{BS}} - 1)(h_{\text{MS}} - 1)f_c/c$ in meters, where c is the speed of light. Here, $h_{\text{BS}} = 25$ m, $h_{\text{MS}} = 1.5$ m, and $f_c = 2$ GHz so that $d_{\text{BP}} = 320$ m:

$$\text{PL}_{\text{basic}}^{(\text{NLoS})}(d) = 9.63 + 39 \log_{10}(d) \text{ in dB}, \quad (3.8)$$

$$\text{PL}_{\text{basic}}^{(\text{LoS})}(d) = \begin{cases} 34.02 + 22 \log_{10}(d) \text{ in dB} & \text{if } 10 \text{ m} < d < d_{\text{BP}} \\ -11.02 + 40 \log_{10}(d) \text{ in dB} & \text{if } d_{\text{BP}} < d < 5000 \text{ m.} \end{cases} \quad (3.9)$$

Correlated Large-Scale Parameters

As a last step before small-scale parameters and channel coefficients can be generated, a set of five correlated large-scale parameters is generated for each link in the scenario. The five large-scale parameters are (1) the delay spread¹ σ_{delay} between the paths of one link, (2 and 3) the angular spreads σ_{AoD} and σ_{AoA} for the departure and arrival angles of all $N_{\text{MPC}} \times M$ rays, (4) the shadow fading, and (5) the Rician K factor specifying the power ratio between the direct LoS ray and the sum power of all other rays. The shadow fading is added to the pathloss to obtain the total large-scale fading of the link. The K factor is used to scale the ray powers in the case of LoS propagation. The other large-scale parameters (σ_{delay} , σ_{AoD} , and σ_{AoA}) are used to parameterize the probability distributions from which the small-scale parameters are obtained. In the following, we present how the large-scale parameters are generated for the example of the UMA scenario with NLoS and LoS propagation. Other scenarios use different parameter sets than those shown in Table 3.4. Note that for NLoS propagation, there are only four large-scale parameters because the Rician K factor is not needed, that is, $K = 0$ because no energy is in a direct LoS ray.

The large-scale parameters are correlated in two ways: First, for each of the five large-scale parameters, a geographic correlation between different links that go to the same base station site is established depending on the geographic proximity of the involved

¹Note that in this chapter and in the context of the IMT-Advanced channel model as defined in [M.2135], the delay spread σ_{delay} large-scale parameter is not used as defined in Section 2.1.1 so that we use a different symbol.

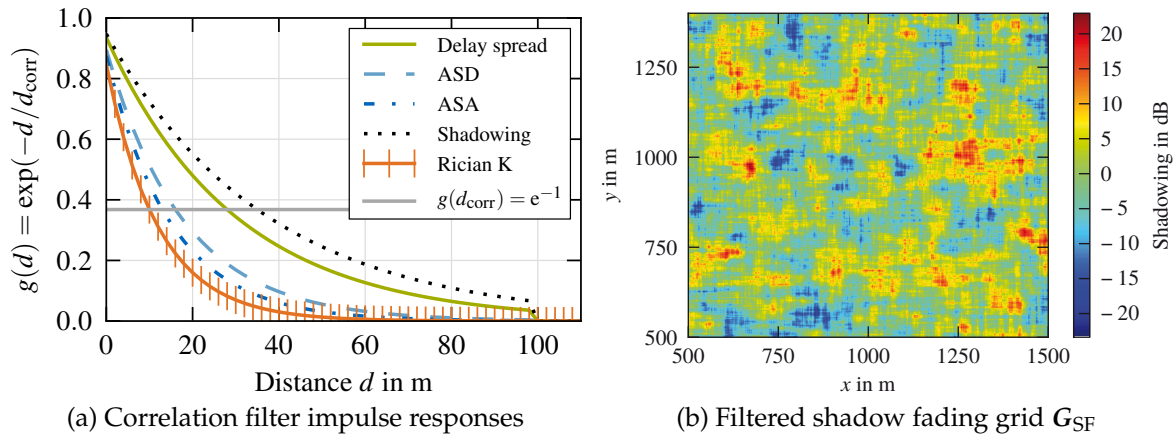


Figure 3.7.: Large-scale parameter filtering for the UMa scenario with LoS

mobiles. This is done separately for each of the five large-scale parameters. For example, the shadowing is identical for links going from one user position to the same site (full intra-site correlation). It is somewhat correlated for links of geographically close users if the links go to the same site and it is independent for links to different sites (no inter-site correlation)². Second, for each link a cross-correlation between the five large-scale parameters is established by means of a correlation matrix.

The ITU has adopted [M.2135] the correlation procedure proposed by the WINNER project [KM+08]. It establishes the correlations in a normal Gaussian domain before transforming the large-scale parameters to their desired distribution: For each of the 19 sites and separately for LoS and NLoS propagation links, five matrices \mathbf{G}_{DS} , \mathbf{G}_{ASD} , \mathbf{G}_{ASA} , \mathbf{G}_{SF} , and $\mathbf{G}_K \in \mathbb{R}^{\dim_x \times \dim_y}$ containing i.i.d. standard normal random variables for each $1 \text{ m} \times 1 \text{ m}$ grid point of the scenario area are created. Note that the dimensions \dim_x and \dim_y of the hexagonal scenarios can be over 1 km so that each matrix contains millions of entries. The geographic correlations inside these matrices are then established by a two-step convolution (e.g., first in the x -dimension and then in the y -dimension) with a one-dimensional *finite impulse response* (FIR) filter whose impulse response $g(d) = \exp(-d/d_{\text{corr}})$ is non-zero only for distances $d = 0 \text{ m} \dots 100 \text{ m}$ between two x - or y -coordinates as plotted in Fig. 3.7a. The parameter d_{corr} specifies the correlation distance, which differs for each combination of scenario and propagation condition, as listed in Table 3.4. Figure 3.7b shows a part of an exemplary filtered shadow fading grid \mathbf{G}_{SF} . After introducing the geographic correlations, for each link a vector

$$\boldsymbol{\zeta}^{(x,y)} = ([\mathbf{G}_{DS}]_{x,y} \quad [\mathbf{G}_{ASD}]_{x,y} \quad [\mathbf{G}_{ASA}]_{x,y} \quad [\mathbf{G}_{SF}]_{x,y} \quad [\mathbf{G}_K]_{x,y})^T \quad (3.10)$$

is created by taking the values corresponding to the (x, y) position of the respective MS from the grids. Now, cross-correlations between the large-scale parameters of each link

²This means that interfering channels to different sites are not correlated in the IMT-Advanced model. The WINNER project states [KM+08] that inter-site correlations were high in one drive test they conducted and low in another one. Because their data was limited, they assumed no correlations for links to different sites although their measurements showed such correlations could exist.

are introduced by a multiplication $s = \sqrt{\mathbf{R}} \cdot \boldsymbol{\zeta}^{(x,y)}$ with a scenario- and propagation-specific correlation matrix with $\sqrt{\mathbf{R}} \cdot \sqrt{\mathbf{R}} = \mathbf{R}$. Equation (3.11), as an example, shows the cross-correlation coefficients for UMa LoS links as specified in report [M.2135]:

$$\begin{matrix} & \text{DS} & \text{ASD} & \text{ASA} & \text{SF} & \text{K} \\ \text{DS} & \left(\begin{array}{ccccc} 1 & 0.4 & 0.8 & -0.4 & -0.4 \\ 0.4 & 1 & 0 & -0.5 & 0 \\ 0.8 & 0 & 1 & -0.5 & -0.2 \\ -0.4 & -0.5 & -0.5 & 1 & 0 \\ -0.4 & 0 & -0.2 & 0 & 1 \end{array} \right) & & & & \\ \text{ASD} & & & & & \\ \text{ASA} & & & & & \\ \text{SF} & & & & & \\ \text{K} & & & & & \end{matrix} = \mathbf{R}^{(\text{UMa,LoS})} \quad (3.11)$$

After element-wise normalization, the vector $\mathbf{s} = (s_{\text{DS}} \ s_{\text{ASD}} \ s_{\text{ASA}} \ s_{\text{SF}} \ s_{\text{K}})^{\text{T}}$ of properly correlated standard normal random variables can now be used to obtain the final large-scale parameters per link. All large-scale parameters are log-normal distributed. For a variable x whose logarithm $\log_{10}(x)$ follows a normal distribution $\mathcal{N}(\mu, \sigma^2)$ with mean μ and standard deviation σ we write $x \sim \log_{10} \mathcal{N}(\mu, \sigma^2)$. The transformations from the normal Gaussian to the log-normal domain are shown in (3.12) to (3.16) with the respective mean and standard deviations given in report [M.2135] (see Table 3.4 for UMa NLoS and LoS values). Note that for NLoS propagation, there is no Rician K large-scale parameter. The grid \mathbf{G}_{K} is thus not needed and \mathbf{R} and $\boldsymbol{\zeta}^{(x,y)}$ have less entries, accordingly:

$$\sigma_{\text{delay}} = 10^{(\mu_{\text{DS}} + s_{\text{DS}}\sigma_{\text{DS}})} \sim \log_{10} \mathcal{N}(\mu_{\text{DS}}, \sigma_{\text{DS}}^2) \quad \text{in s} \quad (3.12)$$

$$\sigma_{\text{AoD}} = 10^{(\mu_{\text{ASD}} + s_{\text{ASD}}\sigma_{\text{ASD}})} \sim \log_{10} \mathcal{N}(\mu_{\text{ASD}}, \sigma_{\text{ASD}}^2) \quad \text{in degree} \quad (3.13)$$

$$\sigma_{\text{AoA}} = 10^{(\mu_{\text{ASA}} + s_{\text{ASA}}\sigma_{\text{ASA}})} \sim \log_{10} \mathcal{N}(\mu_{\text{ASA}}, \sigma_{\text{ASA}}^2) \quad \text{in degree} \quad (3.14)$$

$$\text{Shadowing} = s_{\text{SF}}\sigma_{\text{SF}} \sim \mathcal{N}(0, \sigma_{\text{SF}}^2) \quad \text{in dB} \quad (3.15)$$

$$\text{Rician K} = \mu_{\text{K}} + s_{\text{K}}\sigma_{\text{K}} \sim \mathcal{N}(\mu_{\text{K}}, \sigma_{\text{K}}^2) \quad \text{in dB} \quad (3.16)$$

3.3.2. Initialization of Small-Scale Parameters

After generating the correlated large-scale parameters for each link, the actual small-scale parameters can be determined. The exact process is specified in the ITU report [M.2135]. Here we will only give a rough description: First, the individual path delays τ_n are generated following an exponential distribution according to the delay spread σ_{delay} . The path delays τ_n are normalized by subtracting the smallest delay from all others and then sorted so that $\tau_1 = 0 \leq \tau_2 \leq \dots \leq \tau_{N_{\text{MPC}}}$. Afterwards, the power contribution P_n of each path is derived from the delay distribution. The path powers are normalized to $\sum_{n=1}^{N_{\text{MPC}}} P_n = 1$. Note that at this stage, the path powers are only fractions denoting how the power is distributed among paths. They have to be scaled by the transmit power and the large-scale attenuation effects to get the actual received power in watt. Within each path, the power is distributed uniformly over the 20 rays reflected by the corresponding cluster. Then, arrival and departure angles

$\varphi_{n,m}$ and $\phi_{n,m}$ of each ray m of each cluster n are computed given the angular spread large-scale parameters σ_{AoA} and σ_{AoD} . A wrapped Gaussian distribution (Laplacian for InH) is used to generate arrival and departure angles for the cluster centers. The individual ray angles are then deterministically distributed around the cluster center angle according to a scenario-specific per-cluster angular spread. At this point, arrival and departure rays of the same cluster are randomly coupled to each other. Random initial phases $\Phi_{n,m} \sim \mathcal{U}(-\pi, \pi)$ are finally drawn (separately for all vertical (V) and horizontal (H) polarization combinations: VV, VH, HV, and HH) to model different ray path lengths and phase changes induced by the reflecting material. Before the path coefficients can be computed from these small-scale parameters, the 20 rays of the two strongest clusters (according to P_n) are split into three subclusters each (with 10, 6, and 4 rays each). The delays of the two strongest clusters are spread by fixed offsets of 0 ns, 5 ns, and 10 ns leading to more delay taps in the CIR. This subclustering was introduced in order to model frequency-selective fading for bandwidths of up to 100 MHz [KMH+08].

The IMT-Advanced channel model is said to exhibit three levels of randomness: First, the correlated large-scale parameters are determined from random numbers. Second, the small-scale parameters are drawn randomly from distributions that are parameterized with the correlated large-scale parameters. Third, the initial phases are randomly selected from a uniform distribution. Once all of these parameters have been set up, the channel model behaves deterministically with respect to the model parameters.

3.3.3. Coefficient Generation

The computation of the complex path coefficients $c_{u,s,n}$ at time t follows a sum of sinusoids approach that sums over all rays $m = 1 \dots M$ that belong to a cluster n . Equation (3.17) shows that the summands for each ray consider the horizontal and vertical polarization components of the MS and BS antenna field patterns, the coupling between departure and arrival rays, the phase offsets between the antenna elements in the MS and BS array, and the time-dependent phase shift resulting from the ray's Doppler frequency component $v_{n,m}$:

$$\begin{aligned}
 c_{u,s,n}(t) = \sqrt{P_n} \sum_{m=1}^M & \left[\overbrace{\begin{pmatrix} F_{\text{MS},u}^{(\text{V})}(\varphi_{n,m}) \\ F_{\text{MS},u}^{(\text{H})}(\varphi_{n,m}) \end{pmatrix}^{\text{T}}}^{\text{MS antenna pattern}} \overbrace{\begin{pmatrix} e^{j\Phi_{n,m}^{(\text{VV})}} & \frac{e^{j\Phi_{n,m}^{(\text{VH})}}}{\sqrt{\kappa}} \\ \frac{e^{j\Phi_{n,m}^{(\text{HV})}}}{\sqrt{\kappa}} & e^{j\Phi_{n,m}^{(\text{HH})}} \end{pmatrix}}^{\text{coupling between polarization components with random phases}} \overbrace{\begin{pmatrix} F_{\text{BS},s}^{(\text{V})}(\varphi_{n,m}) \\ F_{\text{BS},s}^{(\text{H})}(\varphi_{n,m}) \end{pmatrix}}^{\text{BS antenna pattern}} \right. \\
 & \cdot \underbrace{\exp(2\pi j d_u \lambda_0^{-1} \sin(\varphi_{n,m}))}_{\text{MS array phase offset element } u} \underbrace{\exp(2\pi j d_s \lambda_0^{-1} \sin(\varphi_{n,m}))}_{\text{BS array phase offset element } s} \underbrace{\exp(2\pi j v_{n,m} t)}_{\text{Doppler phase shift}} \left. \right] \quad (3.17)
 \end{aligned}$$

The computation in (3.17) depends on a number of parameters, which are summarized in Table 3.5. Besides the previously computed small-scale parameters ($P_n, \phi_{n,m}, \varphi_{n,m}$,

Table 3.5.: Spatial channel model parameters

Parameter	Description
$c_{u,s,n}(t)$	complex path coefficient of CIR at time t
s	BS antenna index $s = 1 \dots S$
u	MS antenna index $u = 1 \dots U$
n	path (cluster) index
m	ray index $m = 1 \dots 20$
τ_n	delay of path n
P_n	power fraction of path n
$F_{MS,u}^{(V)}(\varphi_{n,m}), F_{MS,u}^{(H)}(\varphi_{n,m})$	electrical field pattern at MS antenna u and BS antenna s
$F_{BS,s}^{(V)}(\varphi_{n,m}), F_{BS,s}^{(H)}(\varphi_{n,m})$	with vertical (V) and horizontal (H) polarization
$\Phi_{n,m}^{(VV)}, \Phi_{n,m}^{(VH)}, \Phi_{n,m}^{(HV)}, \Phi_{n,m}^{(HH)}$	random initial phases
κ	cross-polarization power ratio
d_s, d_u	distances of antennas s and u to reference antenna element
$\phi_{n,m}$	angle of departure (at the BS) of ray m in cluster n
$\varphi_{n,m}$	angle of arrival (at the MS) of ray m in cluster n
λ_0	wavelength at center frequency
$v_{n,m}$	Doppler frequency component of ray m in cluster n
v and φ_v	mobile's speed vector and direction of travel in the xy-plane

and $\Phi_{n,m}$), a scenario-specific [M.2135] cross-polarization power ratio κ is considered as well as a number of link-specific parameters: The (electrical) field patterns of the BS and MS antennas are derived, for example, from the 3D antenna power gain pattern introduced in (3.3) as $F_{BS,s}^{(V)}(\varphi_{n,m}) = \sqrt{G(\varphi_{n,m}, \theta)}$. If only vertically polarized antennas are assumed, the horizontal field pattern is set to 0 and only the vertical-to-vertical coupling with initial phase $\Phi_{n,m}^{(VV)}$ remains in (3.17). Note that the spatial channel model is two-dimensional only (clusters have no elevation) so that the elevation angle θ for the antenna pattern has to be derived from, for example, the line-of-sight elevation between the BS and MS. For uniform linear arrays at the BS and MS, d_s and d_u denote the distances between the reference antenna element (i.e., $s = u = 1$) and the considered elements s and u , respectively. For arbitrary antenna array layouts, effective distances d_s and d_u have to be computed depending on the arrays' geometries [M.2135]. Finally, the Doppler frequency component $v_{n,m} = \|v\| \cos(\varphi_{n,m} - \varphi_v) / \lambda_0$ is derived from the mobile's speed component along the ray direction with respect to the system center wavelength λ_0 for a mobile traveling in the azimuth direction φ_v . For each link, the path coefficients for all paths $n = 1 \dots N_{\text{MPC}}$ in the CIR are computed according to (3.17). If a link has LoS propagation, a single line-of-sight ray with power $P^{(\text{LOS})} = K / (K + 1)$ is added to the zero-delay coefficient of the first path $c_{u,s,n=1}$ where K is the Rician power ratio large-scale parameter introduced in (3.16). Before adding the LoS ray, all path coefficients are scaled down as $c'_{u,s,n} = \sqrt{1 / (K + 1)} \cdot c_{u,s,n}$ so that the path powers still sum to 1.

The antennas are assumed to be in the far field so that each ray is seen as a planar

wavefront arriving (and departing) with identical angles $\varphi_{n,m}$ at different antenna array elements resulting in spatially correlated fading. Different rays have different Doppler components so that the superposition of rays causes temporal fading on each path. If an FDD system with different center frequencies for DL and UL is modeled, all parameters except λ_0 and the initial phases can be shared between the DL and UL channel model. For large frequency offsets between the DL and UL frequency bands, the DL and UL fading is uncorrelated though.

3.3.4. Computationally Efficient Time Evolution of CIRs and CTFs

Computing the CIR for a time instant t involves an elaborate channel model initialization procedure and then the summation over 20 rays in (3.17) for each of up to 24 (sub)clusters and for each antenna pair—and that is only for a single BS–MS link! Considering that a typical system-level simulation scenario consists of 57 cells with 10 mobiles per cell, we need to compute CIRs and CTFs for $57 \cdot 57 \cdot 10 = 32\,490$ links since we need to model serving as well as interfering links. Thus, for a 4×4 MIMO configuration up to $32\,490 \cdot 4 \cdot 4 \cdot 24 \cdot 20 \approx 250$ million complex ray coefficients have to be computed and summed for each time instant t . Clearly, an efficient implementation is needed to deploy the IMT-Advanced channel model in computer simulations given that system-level simulations should cover a significant time span.

The important insight here is that the specification of the channel model allows to offload the main computation effort from the time evolution phase to the initialization phase. Equation (3.17) can be rewritten to show that only one term depends on the current time t :

$$c_{u,s,n}(t) = \sum_{m=1}^M [a_{u,s,n,m} \exp(jb_{n,m}t)] \quad (3.18)$$

This means that at the beginning of a drop, after the large-scale correlation and small-scale initialization processes have been performed, we can compute the constants $a_{u,s,n,m}$ and $b_{n,m}$. As the pathloss (PL) and shadow fading (SF) are fixed for the duration of the drop, we can incorporate these large-scale effects by scaling the $a_{u,s,n,m}$ coefficients accordingly with $\sqrt{\text{PL} \cdot \text{SF}}$. For the above example, 250 million complex-valued constants $a_{u,s,n,m}$ are still needed though. But now an efficient computer implementation as, for example, in `IMTaphy` is possible: We can stack the constants $a_{u,s,n,m}$ and $b_{n,m}$ for all links, antenna pairs, paths, and rays into huge vectors \mathbf{a} and \mathbf{b} . Then for each new time evolution step t , we scale the vector \mathbf{b} by t . After that, we feed each element of the scaled vector \mathbf{b} into the `cis` function defined as $\text{cis}(x) := \exp(jx)$. Both the scaling and the `cis`(x) function are available in vector math libraries that distribute the workload over different processors and utilize the vector extensions of modern processors. The element-wise multiplication of the `cis`(x) output with the vector \mathbf{a} , followed by a summation over $M = 20$ rays, can then be executed as multiple dot products of two length- M vectors for each link, antenna pair, and path.

The last step in the computation of the CTF channel matrices is the DFT transforming the time-domain CIR coefficients per path to the frequency domain for each subcarrier

of interest. In general, FFT algorithms are the preferred way to compute DFTs. However, they typically operate on a uniformly spaced grid with as many bins in the time domain as there are frequency bins. In contrast, we only have up to 24 coefficients in the time domain at arbitrary delays τ_n that need to be transformed to, for example, 1200 subcarriers in a 20 MHz system. Instead of sampling the path coefficients to a uniform time grid filled with mostly zeros and performing an FFT, it turns out that a direct DFT can be significantly faster. With a suitable array layout that avoids copying intermediate results in the computer's memory, the CIR coefficients can be regarded as a sequence of $\mathbf{Z} \in \mathbb{C}^{(U \cdot S) \times N_{\text{MPC}}}$ CIR matrices for all links, where U is the number of MS antennas and S is the number of BS antennas, see Table 3.5. By matrix multiplications, the frequency domain CTF channel matrices $\mathbf{H}' = \mathbf{Z}\mathbf{F}$ with $\mathbf{H}' \in \mathbb{C}^{(U \cdot S) \times F}$ can be computed from the Fourier coefficients $\mathbf{F} \in \mathbb{C}^{N_{\text{MPC}} \times F}$ with $[\mathbf{F}]_{n, f_{\text{idx}}} = \exp(-2\pi j f_c(f_{\text{idx}}) \tau_n)$, where $f_c(f_{\text{idx}})$ is the center frequency of subcarrier $f_{\text{idx}} = 1 \dots F$. Using a different memory indexing, the coefficients of the CTF channel matrices $\mathbf{H}' \in \mathbb{C}^{(U \cdot S) \times F}$ can be accessed as a series of frequency domain channel matrices $\mathbf{H}(f_{\text{idx}}) \in \mathbb{C}^{U \times S}$ organized in a conventional Rx-by-Tx layout (here for the DL).

High-efficiency implementations of all of the discussed vector and matrix operations are included in standard vector math libraries. Since such libraries also exploit the vector extensions of modern CPUs, using single-precision floating point numbers results in twice the computational throughput with half the memory consumption compared to double-precision numbers. The slightly reduced numeric precision can be tolerated because it results in a small random error affecting an anyway stochastically generated channel. On a 2011 off-the-shelf desktop computer with a single processor (four cores), the `IMTphy` implementation of the IMT-Advanced channel model is able to perform a time evolution step for a 32 490 links UMa scenario with a 4×4 MIMO configuration in less than one second. About 10 % of that time is spent on the DFT to derive the CTF on 100 frequency bins. Considering that during this one second about 400 MByte of CTF matrices are updated in the computer's memory, an on-the-fly computation approach like the one taken by `IMTphy` probably outperforms loading pre-computed channel data from an external memory. When both downlink and uplink are simulated simultaneously in an FDD system, almost twice the CPU time is needed because equations (3.17) and (3.18) depend on the center frequency.

3.4. Channel Model Calibration

Different companies and organizations involved in the 3GPP standardization or in the IMT-Advanced evaluation process have implemented the IMT-Advanced evaluation guidelines together with the IMT-Advanced channel model in their own simulation tools. Due to a myriad of parameters and because of the complexity of system-level simulators and channel model implementations, simulation results published by different sources usually differ, and sometimes significantly so. There are multiple reasons for the divergence of simulation results even for a fully standardized system like LTE Rel-8. First, system specifications have many parameters that can be freely chosen and

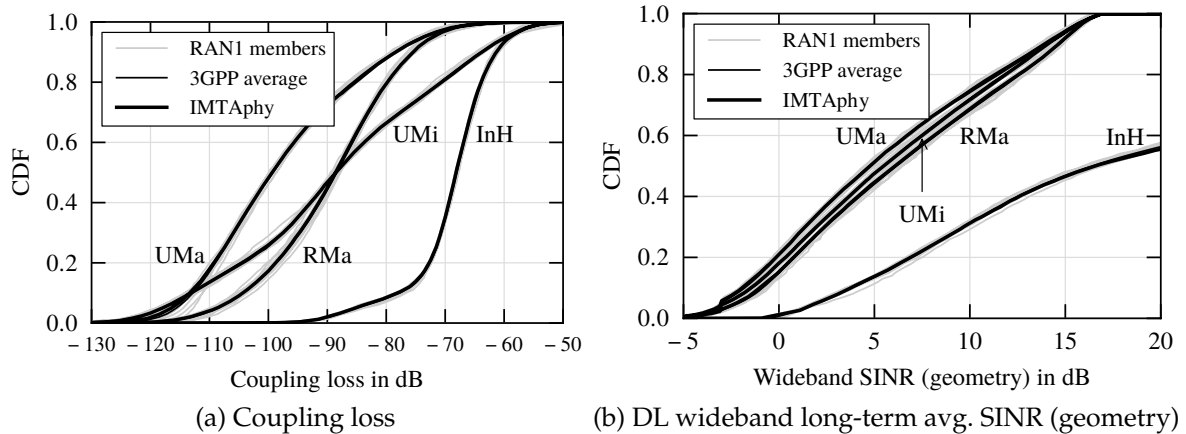


Figure 3.8.: Large-scale calibration of IMTAphy against 3GPP results [R1-092019]

different organizations might choose different settings. Second, certain algorithms like the scheduling algorithm are not standardized and some organizations might have better solutions than others. Third, even detailed models and guidelines sometimes leave room for interpretation. For example, the ITU-R guidelines [M.2135] published in 2008 had to be clarified in 2009 [Fin09a] to avoid different interpretations. Fourth and finally, simulation software—like all software—is usually not free from implementation errors. To avoid the two latter error sources, the 3GPP as well as some of the independent evaluation groups like the WINNER+ project conducted calibration campaigns to align the simulation results of their member organizations. They aimed at excluding the first two mentioned reasons by a detailed specification of simulation parameters and algorithm assumptions. During the course of the calibration campaigns, a common understanding of how to interpret the guidelines was reached as well. A system-level simulation relies on a number of random inputs. Thus, not specific results but rather statistical evaluations like the distribution of certain metrics can be used for the calibration process.

In this section, we will present the channel model calibration efforts by the 3GPP and some IMT-Advanced evaluation groups and show that our implementation in IMTAphy complies to the IMT-Advanced modeling guidelines. In addition, we will introduce some further metrics that can be used to verify the generated channel matrices. In Chapter 4 we will present more aspects of the 3GPP's calibration campaign that are specific to LTE.

3.4.1. Large-Scale Calibration Metrics

The two large-scale calibration metrics presented here are the *coupling loss* CDF and the *geometry* CDF that are evaluated over all links. The *coupling loss* is the sum of all the losses between an MS and its serving BS without considering the fast-fading: the antenna feeder loss, the influence of BS and MS antenna patterns, the pathloss on the link, potential O2I or *outdoor-to-vehicle* (O2V) penetration losses, and the shadowing.

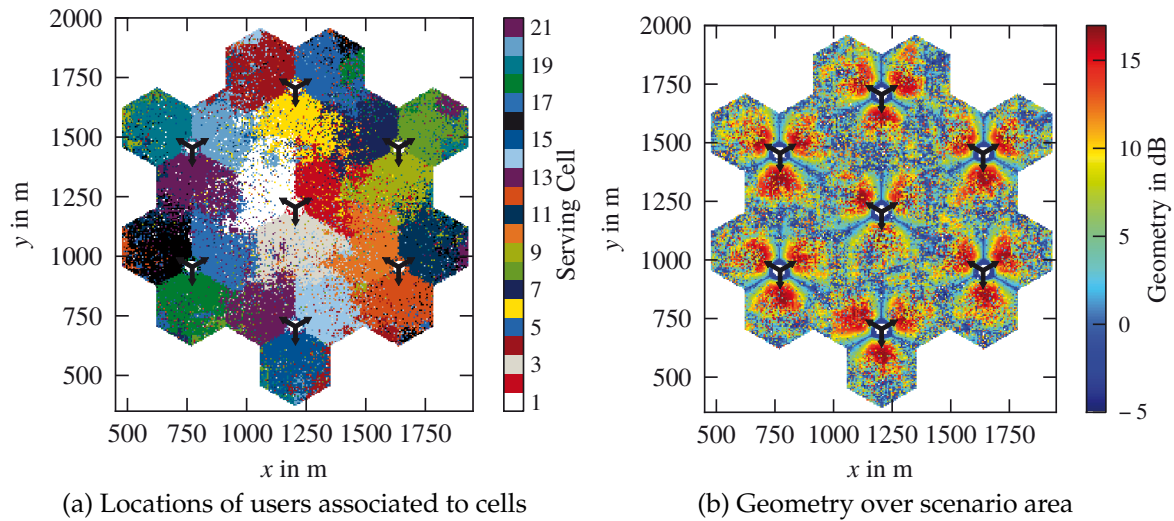



Figure 3.9.: Visualization of a small UMA scenario with 7 sites and 21 cells

Besides testing the implementation of the mentioned components, the CDF of the coupling loss, shown in Fig. 3.8a, also helps to identify problems with the LoS/NLoS classification, the user placement (i.e., wrap-around and minimum distances), and the association process between BSs and MSs (handover hysteresis). Figure 3.9a illustrates to which serving base station users are associated. Each pixel in the figure represents a different MS located at that position. Due to shadowing effects, the associations do not perfectly follow the hexagonal cell deployment. Also, the wraparound modeling, which links cells at opposite sides of the scenario, is clearly visible in Fig. 3.9a because it shows some users in scenario-edge cells associated to base stations serving cells on the other side of the scenario.

The *geometry* is the per-user wideband and long-term downlink linear SINR average in a full reuse scenario with round-robin scheduling. Figure 3.9b shows its distribution over the considered small scenario area. The antenna sector patterns and the spatially-correlated shadowing can be identified, as well as a random O2V penetration loss component, which creates a noise-like effect. The geometry CDF shown in Fig. 3.8b helps to reveal correlation issues between links, for example, if different propagation conditions are applied to MS–BS links going to the same site. It can, however, hide modeling problems with one of the mentioned coupling loss factors if the losses on the serving and interfering links are off by the same factor. Figure 3.8 shows the results obtained with IMT_Aphy (averaged over multiple random drops) compared to the results obtained by more than a dozen 3GPP *radio access network working group 1* (RAN1) members who participated in the 3GPP calibration campaign [R1-092019]. As can be seen, there are only small deviations between the curves of the 3GPP RAN1 members. From the RAN1 members' individual results an average curve is computed (see also TR [36.814]), which is almost identical to and thus concealed by the thicker IMT_Aphy curve. A technical document [R1-091482] presents CDFs for LoS and NLoS propagation conditions separately, which can be very helpful for debugging purposes.



(a) IMTAphy curves compared to reference results [WIN10; Chi] in the UMa scenario (b) IMTAphy delay spread scenario comparison

Figure 3.10.: CDFs of the RMS delay spreads computed according to TR [25.996]

3.4.2. Small-Scale Calibration Metrics

To test the small-scale parameters that are used to generate the channel coefficients, two metrics defined in TR [25.996] are commonly used: The CDF (again over all links) of the weighted root mean square (RMS) delay spread $\sigma_{\tau}^{(\text{RMS})}$ as defined in (3.19) and the CDF of the weighted RMS circular angular spread. For each link, the weighted RMS delay spread $\sigma_{\tau}^{(\text{RMS})}$ is computed from the delay of each path τ_n weighted by that path's power P_n :

$$\sigma_{\tau}^{(\text{RMS})} = \sqrt{\frac{\sum_{n=1}^{N_{\text{MPC}}} (\tau_n - \mu_{\tau})^2 P_n}{\sum_{n=1}^{N_{\text{MPC}}} P_n}} \quad \text{with} \quad \mu_{\tau} = \frac{\sum_{n=1}^{N_{\text{MPC}}} \tau_n P_n}{\sum_{n=1}^{N_{\text{MPC}}} P_n}. \quad (3.19)$$

Note that for LoS propagation links, the first cluster's power also contains the power of the line-of-sight ray.

Figure 3.10a shows the CDFs for the delay spreads in the UMa scenario in comparison to reference results provided by WINNER+ project members [WIN10] and two members (China Mobile Communications Corporation (CMCC) and China Academy of Telecommunication Research (CATR)) of the Chinese IMT-Advanced evaluation group [Chi]. In addition, Fig. 3.10b illustrates the delay spreads derived with IMTAphy for NLoS and LoS propagation in the four IMT-Advanced deployment scenarios.

The RMS angular spread $\sigma_{\text{ASA}}^{(\text{RMS})}$ of arrival angles is computed in a similar way. The definition, however, is a bit more technical because a minimization over all possible rotations of the coordinate system [Nor03] and a normalization of angles between intermediate steps [WIN10] has to be performed to obtain a well-defined metric. For that reason, we omit the definition here and just present the results. Figure 3.11 shows the CDFs of the RMS angular spreads for arrival and departure angles in the UMa scenario generated by IMTAphy compared with the results by members of the WINNER+ and

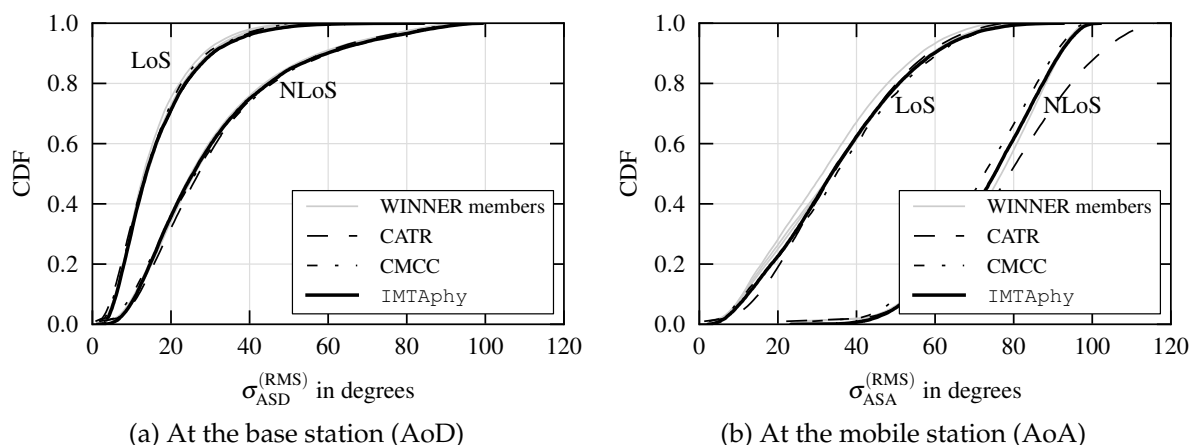


Figure 3.11.: CDFs of the RMS angular spreads computed according to TR [25.996] and reference curves [WIN10; Chi] in the UMa scenario

Chinese IMT-Advanced evaluation groups. It can be seen that the rooftop BS antennas experience a much smaller angular spread than the mobile users who are typically surrounded by scatterers. Both delay and angular spreads for LoS links are significantly smaller because most power is concentrated in the LoS ray of the first cluster. Again, the IMTAphy calibration curves and the curves provided by most organizations match very well. This calibration approach not only allows to test if the delay and angle small-scale parameters generated by an implementation of the channel model match RMS delay and angular spread reference CDFs. It also allows to verify the cluster power small-scale parameters and the Rician K factor distributions because these two additionally influence the delay and angular spread computations.

3.4.3. CIR and CTF Calibrations

The small-scale calibrations presented above only cover the inputs to the channel matrix computation in the generic model. In order to validate the resulting outputs, further tests are desirable. One way to check if the time evolution and the speeds are handled correctly, is to evaluate the temporal autocorrelation properties of a single path of the generated CIR [NST+07]. Figure 3.12a shows that the IMTAphy and the WINNER implementations [Fin09b] of the IMT-Advanced channel model, as well as the older 3GPP SCM channel model [SDS+05], exhibit a very similar temporal autocorrelation. The generated fading also matches the expected theoretic Rayleigh fading autocorrelation function, which is a 0-th order Bessel function of the first kind.

A common way [CWK+09; NST+07] to characterize the correlation properties of the frequency-domain channel matrices \mathbf{H} produced by a channel model, is to consider the distribution of the resulting MIMO capacity $C_{No\ CSIT}$ (see (2.18) on page 16). Figure 3.12b shows the capacity CCDFs comparing IMTAphy, the WINNER implementation [Fin09b], and the SCM channel model [SDS+05] against the capacity distribution of purely i.i.d. random (circularly symmetric complex Gaussian) channel matrices.

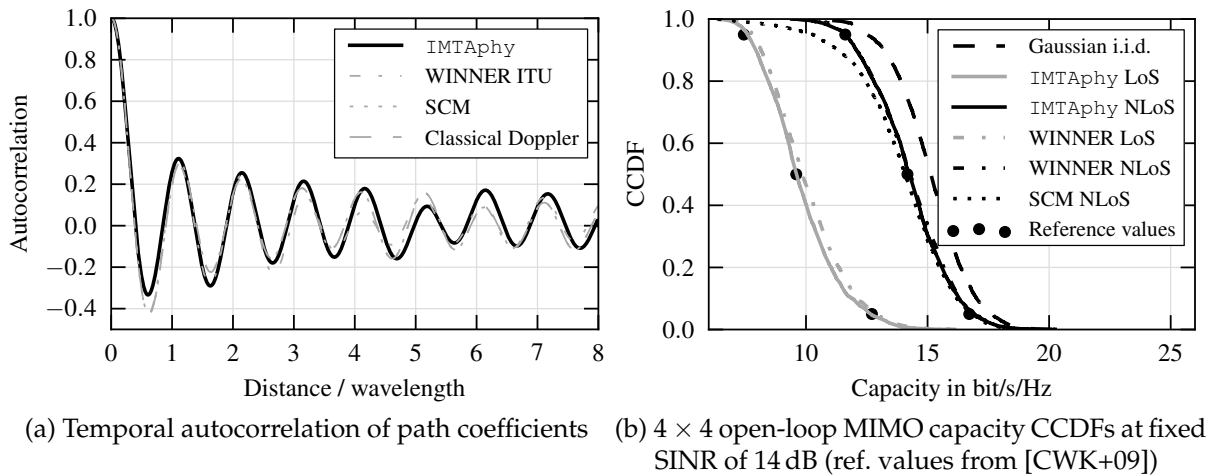


Figure 3.12.: Comparing channel model implementations in the UMa scenario

For all models, the values are generated assuming a 4×4 uniform linear array (ULA) MIMO configuration with 10λ element spacing at the BS and 0.5λ at the MS for an SINR of 14 dB in the UMa scenario. In addition, tabulated values for the 5%, 50%, and 95% CDF levels from Chong et al. [CWK+09] are plotted as a further reference.

One can see that the `IMTAphy` and `WINNER` implementations of the IMT-Advanced channel model match very well. The `SCM` model, for which only NLoS propagation is available, shows a slightly different performance, which is to be expected because the modeling assumptions are not identical to the IMT-Advanced channel model. Obviously, the correlations considered by the different channel models, which are based on measurement campaigns, yield lower capacities than in the Gaussian case, especially for LoS conditions. This observation underlines that relying on a simplistic MIMO channel model that does not account for the channel correlations observed in practice, would lead to a significant overestimation of channel capacities and spectral efficiencies. It also shows that careful modeling of the LoS/NLoS propagation probabilities in the system is necessary for deriving meaningful expected spectral efficiencies.

3.5. Summary and Contributions

The channel model is a major component of the `IMTAphy` simulator, which we developed for the performance evaluation of LTE and the further studies presented in this thesis. As the name `IMTAphy` indicates, it follows the ITU guidelines for the evaluation of IMT-Advanced systems. Making the source code available on the Internet contributes to the idea of reproducible research, and strictly adhering to the de-facto standard set by the ITU makes our results comparable with those of others.

This chapter provided an overview of the IMT-Advanced deployment scenarios, the antenna modeling, and the IMT-Advanced channel models. It further discussed how our `IMTAphy` implementation was calibrated against reference results.

The IMT-Advanced indoor hotspot scenario, as well as the outdoor scenarios for urban micro and macro cells (UMi and UMa), and the rural macro cell (RMa) scenario are the four scenarios we use for our performance evaluation. We focus on the outdoor scenarios because their size (57 cells) and the wraparound model, which we also described, are supposed to offer a realistic model of actual large-scale cellular systems. The scenarios differ with respect to the channel model parameters and deployment characteristics. One main difference in deployment characteristics is the assumed carrier frequency and user mobility velocity, which determine the coherence time.

For the antenna modeling, a combination of ITU and 3GPP simulation guidelines is followed. The azimuth and elevation antenna gain pattern of the sector antennas used in the outdoor scenarios are defined in the IMT-Advanced guidelines. BS antenna array configurations and polarization modeling for the array elements is not provided by the ITU and is thus sourced from 3GPP guidelines.

The channel models distinguish between scenario-specific large-scale and small-scale propagation effects. The dominant large-scale effect is the distance-based pathloss. It depends on whether a BS–MS link exhibits LoS or NLoS propagation, which in turn is modeled using a scenario-specific probability function. We contribute an evaluation of LoS conditions in the IMT-Advanced scenarios showing that all scenarios exhibit a high percentage of serving links with LoS propagation. Besides the LoS versus NLoS classification, four so-called large-scale parameters are derived in a way that ensures that geographically-close users experience correlated conditions. The first is the shadow fading. The others are employed to parametrize the probability distributions used to obtain the delays as well as the arrival and departure angles of multipath reflections. Instead of explicitly modeling the scatterers (e.g. buildings) that cause the multipath propagation environment, the geometry-based stochastic channel model uses these suitably created random values to define the channel impulse response per link.

A significant contribution of this work is the formulation of the channel impulse response computation and of the Fourier transformation in an efficient way. This enables `IMTAphy` to simulate large-scale systems with high MIMO orders in a time-efficient manner by exploiting time-invariant factors in the CIR definition and optimizing the implementation for modern multi-core CPUs featuring vector extensions. Besides giving an overview of modeling guidelines, another important contribution presented in this chapter is the presentation of different calibration steps that were performed to ensure that the `IMTAphy` channel model implementation is correct and delivers useful results. To this end, various large-scale and small-scale parameter distributions were compared to reference results provided, among others, by industry-leading companies participating in the 3GPP standardization process.

To the best of the author’s knowledge and compared to other simulation tools available to academia, the channel model implementation in `IMTAphy` is the only one to combine the IMT-Advanced deployment scenario modeling with the spatial channel models, to offer an efficient implementation enabling fast on-the-fly channel generation of MIMO channels, and to be fully calibrated against 3GPP reference results.

4. Link-to-System Modeling for LTE System-Level Simulations

After having discussed the IMT-Advanced channel model in the previous chapter, we now turn our attention to the link-to-system model, which is the next higher building block in the system-level modeling overview shown in Fig. 2.14 on page 49. It models the performance of the physical layer as depicted in Fig. 2.4 on page 25 with a computational complexity that makes it feasible to simulate thousands of links in a system-level simulation. Even though link-to-system modeling techniques are a standard feature of system-level simulations, they are not specified in detail for IMT-Advanced evaluations. In its guidelines, the ITU merely mentions their use and requires them to be “sound in principle” [M.2135].

In this chapter, we first motivate the need for link-to-system modeling techniques and present various alternatives from the literature. After that, we describe a standard model [BAS+05; EMD] that is also implemented in `IMTAphy` and used to generate all system-level simulation results in this thesis. As a component of the link-to-system model, we discuss in Section 4.2 how the *link quality* of a transmission can be characterized by means of a vector of SINR values. Based on this input, we show in Section 4.4 how the throughput-relevant block error probability can be derived to characterize the *link performance*. We conclude the chapter with calibration results focusing on spectral efficiency that validate the link-to-system model implemented in `IMTAphy` against 3GPP calibration reference results.

While the thesis focuses on the downlink, we will briefly touch uplink modeling and calibration topics in this chapter because `IMTAphy` supports the uplink as well.

4.1. Overview of Link-to-System Modeling for System-Level Simulations

Before giving an overview of link-to-system models in the literature, and as a motivation for the link-to-system model chosen in `IMTAphy`, we first characterize *link-level* simulations and argue that their detailed modeling and corresponding computational complexity prohibit their use in system-level simulations of multi-cell systems.

4.1.1. Link-Level Simulations

Link and system-level simulations differ in the simulation fidelity and complexity. In link-level simulations, the complex baseband transmit signal $s(t)$ is generated exactly

as it would result from the transmitter. This means that the transmission chain shown in Fig. 2.8 is modeled on the bit and resource element level. The transmit signal $s(t)$ is convolved with the channel impulse response $c(t)$ to obtain the received signal $y(t) = c(t) \star s(t) + n(t)$ with \star denoting the convolution operator. Most often, only noise $n(t)$ but no interfering transmissions are considered (single cell). This received signal $y(t)$ is fed into the receiver implementation, which then performs all the processing steps that would happen in an actual receiver. For example, it removes the cyclic prefix, performs an FFT, channel estimation, MIMO detection, and Turbo decoding to obtain all bits of the transmitted transport block. The transport block is received correctly if the decoded bits match the transmitted ones. The entire transport block is discarded (or retained for a future HARQ retransmission) if at least one unrecoverable bit error leads to a failing CRC check.

Conducting a system-level simulation, that is, a simulation in a multi-cell scenario with thousands of links, is generally considered as too complex at this level of detail. The transmit signal would have to be created for all serving links in all cells, and the convolution of the transmit signals with all serving and interfering channel realizations would have to be computed followed by the above described processing in the receiver chain. As we will see in the following, link-to-system models help to accurately characterize the block error probability based on the channel realizations and resource allocation decisions in all cells without the complexity associated with link-level simulations. For a given channel realization and transmit strategy (i.e., precoding, MCS and transport block size), it is the expected BLER that characterizes the throughput some receiver implementation can achieve.

Nevertheless, link-level simulations are still a necessary and useful tool. They can deliver more exact results especially in situations where the influence of inter-cell interference, which is only modeled in system-level simulations, can be neglected. On the one hand, they are needed to generate lookup tables used in link-to-system models. On the other hand, they allow to evaluate the performance of important transceiver aspects such as synchronization, cell search, random access procedures, detection and decoding of control channels, etc., which are usually not in the scope of system-level simulations.

4.1.2. Link Performance Models in the Literature

There are a number of link performance modeling approaches that have been used in the literature. They differ with respect to their computational complexity and regarding the amount of transmission schemes that can be modeled. In the following, we will discuss a few examples before we present the most commonly used model that is also employed by IMT_Aphy. A common feature of all presented models is that they rely on the SINRs with which a transmission is received to model the *link quality*.

The simplest link performance models are intended to be used together with simple channel models that just model pathloss and shadowing on SISO links. Based on these large-scale propagation properties, the SINR of a transmission is determined (link

quality). The link performance is then determined by mapping the SINR to a scaled-down Shannon capacity that is capped at the maximum supported link efficiency. For example, a 3GPP technical report [36.942, Sec. A.1] suggests to model the link spectral efficiency in the LTE downlink for a given SINR between a minimum SINR of -10 dB and a maximum link spectral efficiency of 4.4 bit/s/Hz as

$$\text{link spectral efficiency} = 0.6 \cdot \log(1 + \text{SINR}) \text{ in bit/s/Hz.} \quad (4.1)$$

There are several drawbacks associated with this simplistic modeling. On the one hand, it is not suited to model varying SINRs per transport block due to time and frequency-selective fading or interference. A simple simulation model could thus only model the large-scale fading effects due to pathloss and shadowing and assume that the fast fading “averages out” in the long run. Also, the scaled Shannon mapping is mostly suited for SISO transmissions since MIMO transmission schemes add further challenges. While, for example, the Tx-side beamforming or the Rx-side combining gain could in most cases be accounted for in form of fixed assumed gains (e.g., 3 dB gain if two Rx antennas instead of one are used, or via performance gains published in the literature [CGE03]), the mutual interference between spatial multiplexing layers is hard to quantify without modeling the MIMO channel and receiver in detail. In addition, also the interference impact on the SINR depends on the type of receiver and on the spatial signatures of the interfering signals.

Besides the difficulty of determining the link quality in the form of SINRs in fading MIMO environments, the simple mapping to a throughput or spectral efficiency value using the scaled Shannon mapping has serious limitations in that it neglects the challenges of link adaptation in actual systems. Even if the SINR is accurate, the associated link spectral efficiency is only achieved if the base station selects an appropriate modulation and coding scheme. If the CSI feedback indicates too good SINRs, the chosen MCS might be too aggressive leading to a block error and zero achieved throughput. On the other hand, a too pessimistic SINR assumption would lead to the selection of a too conservative MCS that can be received correctly but falls short of the possible link spectral efficiency. We thus highlight that a performance model needs to account for the link adaption in order to capture the performance impact of this important RRM scheme.

Further challenges to simple link-to-system models based on non-fading channel models arise from the use of channel-adaptive scheduling algorithms like the proportional fair scheduler that was introduced in Section 2.2.3. Since the PF scheduler aims at assigning parts of the radio resources that offer favorable channel conditions, it alters the SINR distribution and allows to obtain an effective SINR gain, if accurate channel feedback is used for the scheduling and link adaption within the channel’s coherence time. One way to model the performance impact of this RRM scheme without the need for a fast-fading channel model was suggested in [Ell09b] by introducing an effective SINR gain that depends, among other factors, on the pre-scheduling long-time SINR and the number of co-scheduled users.

Based on these observations, we conclude that in order to capture the performance of modern transmission schemes used in LTE, accurate link-to-system models need

to operate on fast-fading MIMO channels as those generated by the IMT-Advanced channel models. One approach suggested by Pauli et al. [PVB+09] is to generate both the fading MIMO channel realizations, as well as the resulting performance indicators of MIMO transceivers, offline for later use in system-level simulations. The authors suggest to condense the impact of the channel and the MIMO scheme into a scalar effective gain and an inter-layer interference term in the case of spatial multiplexing. By storing these time-varying factors in so-called fading files, they do not need a channel model implementation for online computation of the channel. As we have discussed in Section 3.3.4, however, with an efficient channel model implementation as in `IMTAphy`, it is possible to generate channels online faster than they could be read from disk. The approach of the authors in [PVB+09] further eliminates, for example, the need for online MIMO filter computations. But for the same reason, it cannot model the impact of channel estimation errors and interference on the MIMO linear receiver performance. In addition, due to disk space limitations, the authors suggest to take the fading information for multiple links from different positions in the same file. That way, no realistic mixes of MIMO channel correlation and time and frequency-selective fading behavior are available.

4.1.3. The Standard Link-to-System Modeling Concept

Figure 4.1 illustrates the current standard link-to-system modeling concept that was introduced by Brüninghaus et al. [BAS+05]. It is recommended for system-level simulations in the 802.16m guidelines [EMD] and is implemented in our `IMTAphy` simulator, as well as in the openWNS [BMKS09], LTE-Sim [PGB+11], LENA [BMRN11], and in the TU Vienna [Iku13] system-level simulators. It can be used independently of the MIMO channel implementation so that, for example, MIMO channels generated with the IMT-Advanced model or even raytracing-based channels could be employed.

Given a transmit strategy, that is, the allocated PRBs, chosen precoders and the modulation and coding scheme, the expected throughput that can be achieved on the link depends on the channel realizations and interference conditions as well as on the performance of the receiver implementation. The presented link-to-system model thus consists of a *link quality model* that captures the SINR conditions on the serving and interfering links and accounts for the impact of linear receive filters. In the downlink protocol stack illustration in Fig. 2.4 on page 25, the link quality model comprises all the stages in the transceiver chain from the modulation mapping on the transmitter side to the output of the inner receiver. The performance of the Turbo decoder in the outer receiver in Fig. 2.4 is then modeled by the *link performance model*, which depends on the modulation and coding scheme in addition to the inputs it receives from the inner receiver. Instead of log-likelihood ratios as in an actual receiver, the link quality input to the performance model consists of a vector of SINRs [BAS+05; EMD].

In the example in Fig. 4.1, a user has been allocated a transport block spanning two layers and two PRBs colored in blue. Taking the time- and frequency-selective channel matrix $\mathbf{H}(t, f)$, the chosen precoder $\mathbf{P}(t, f)$, as well as the interference covariance

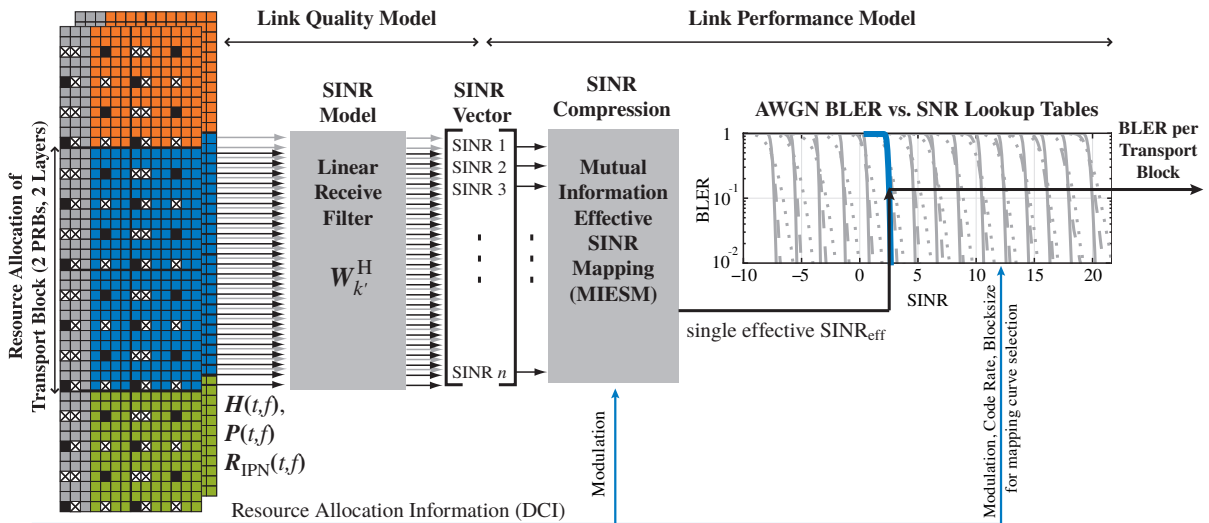


Figure 4.1.: Overview of link-to-system modeling methodology determining block error ratio for PDSCH allocation of single transport block (blue) on two PRBs and two layers [Iku13; BAS+05]

$R_{\text{IPN}}(t, f)$ on these physical resources into account, the link quality model computes a vector of post-MIMO-equalization SINRs (see also Section 4.2). SINRs are computed for both layers and, depending on the modeling granularity, for the allocated PRBs or subcarriers (two PRBs or 24 subcarriers in the example in Fig. 4.1). The individual SINRs can differ due to frequency-selective fading and interference, different precoding between the layers, and due to temporal fading if multiple symbols during one TTI are to be modeled. As the different SINRs all belong to the same transport block that is transmitted with the same modulation and code rate over the whole physical resource allocation, the vector of SINRs is compressed into a single representative *effective SINR*. In IMT_{phy}, the compression is done using a so-called *mutual information effective SINR mapping* (MIESM), see Section 4.4.1. Using this single SINR value per transport block, the block error ratio is taken from a pre-computed table (plotted in Section 4.4.2) according to the used modulation, code rate, and code block length.

The link quality model accounts for the capacity losses a) to d) in Table 2.2 on page 43, which are due to the use of codebook-based precoders and linear receivers. The remaining transceiver losses, that is, e) to g) in Table 2.2 are accounted for by the link performance model. The effect of BICM with QPSK, 16-QAM, and 64-QAM is at the heart of the mutual information-based SINR mapping. In addition, the BLER lookup tables capture the performance of a Turbo encoder and decoder algorithm under realistic code block length and code rate conditions. The performance of the Turbo decoder is implementation-specific and, for example, depends on the number of iterations so that a representative algorithm is evaluated offline to generate lookup tables. Finally, the signaling and channel estimation overhead in the LTE resource grid discussed in Section 2.3.2 is also considered in our simulations. For each modulation and coding scheme, the standard [36.213, Tab. 7.1.7.2.1-1] lists the size of the (uncoded) transport

block in bits depending on the number of allocated PRBs. The code rate (i.e., the ratio of uncoded payload bits to the number of bits that fit into the allocated resources) results from the number of available REs, which depends on the RE overhead given in Table 2.3. It is the task of the link adaptation RRM method in the base station to select the MCS index that yields the desired code rate given the overhead situation and the estimated channel conditions. The transport block and resource grid size definitions account for the overhead due to the OFDM cyclic prefix. The frequency guard bands are also respected because only the net bandwidth (e.g., in a 20 MHz system a total of 100 PRBs encompassing 19 MHz) is made available to the simulation. Altogether, the link-to-system models in our system-level simulations consider all types of losses discussed in Sections 2.3.1 to 2.3.3. The simulation results thus indicate the net throughput that is available to the layer 3 of the protocol stack.

Note that the link quality model presented here only allows to model the performance of linear receivers as introduced in Section 2.2.2. For the reception of spatial multiplexing transmissions, more advanced MLD-like non-linear receivers (e.g., sphere decoders) could be used that perform a joint detection of the symbols transmitted on all layers, see the discussion of maximum likelihood detection in Section 2.2.2. In contrast to the situation with linear receivers, which perform a symbol-by-symbol detection independently per layer, it is not-straightforward to assign link quality measures (e.g. SINRs) to the individual layers that are detected jointly. This not only makes the link-to-system modeling of MLD-like receivers difficult, it also affects the feedback-based link adaptation performance in actual systems because the MS has to estimate its future ML detection performance given the current channel and interference conditions without having a concrete spatial multiplexing transmission to receive.

Suggestions for incorporating MLD performance modeling into link-to-system models have already been presented in the 802.16m evaluation methodology document [EMD]. In contrast to the performance models for linear receivers presented in this chapter, however, these initial MLD models have not found widespread adaptation and can thus not serve as the basis for a representative performance evaluation. Recently, new abstraction techniques for ML-detection have been proposed [MLKL12] that promise to deliver accurate results with lower computational complexity compared to the earlier methods suggested in [EMD]. The authors [MLKL12] suggest an approach that is compatible with the general link-to-system modeling concept presented here by only modifying the link quality model part to support MLD receivers. In the proposed model, a per-layer SINR is estimated by bounding the MLD performance from below by the MMSE performance and from above by the ideal inter-layer-interference-free performance. Since this newly proposed scheme, like the previous ones in [EMD], relies on sets of tuning parameters that have to be adapted for the used modulation and coding schemes, an MLD-compatible link quality model is not as well-defined as the link quality model for linear receivers that we will present in the following. Consequently, for simulator calibration purposes and as a baseline LTE-Advanced receiver performance assumption, only linear receivers are considered, for example, by the 3GPP [36.814]. Thus, in *IMT_Aphy* and for the performance evaluation of LTE-Advanced in the remainder of this thesis, we only use linear receivers.

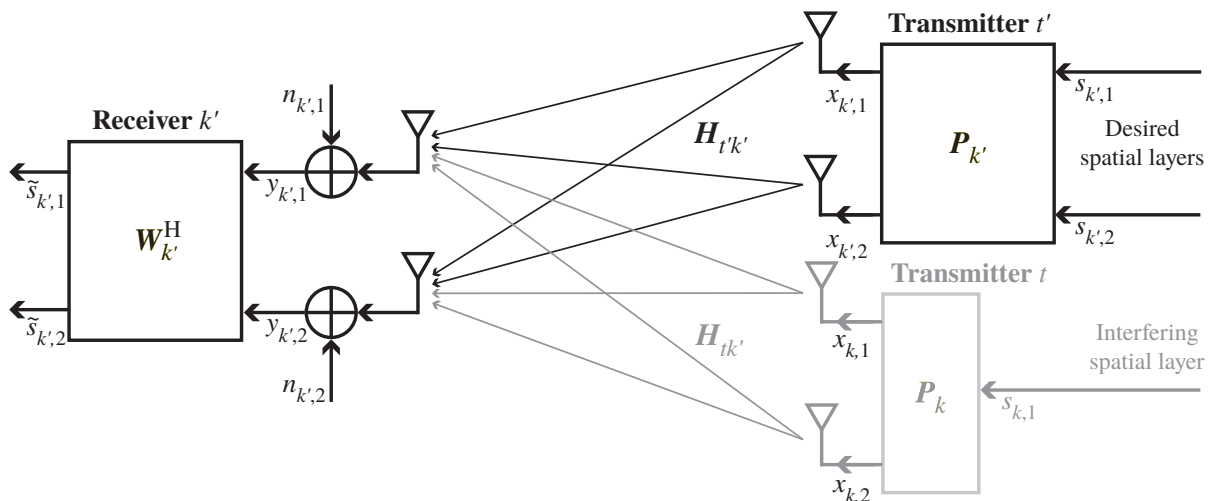


Figure 4.2.: Receiver k' employing a filter matrix $W_{k'}^H$ to receive two spatial layers $s_{k',1}$ and $s_{k',2}$ over a 2×2 MIMO channel $H_{t'k'}$ from a desired transmitter t' using a precoding matrix $P_{k'}$ with interference coming from an undesired transmitter t using a precoder P_k to transmit to some other user $k \neq k'$

4.2. Modeling of MIMO Linear Receivers and Precoders

In this section, we present the link quality model used in `IMTaphy` that allows to compute the post-equalization SINRs for a MIMO system with linear receivers and precoders. We consider a MIMO system as depicted in Fig. 4.2 and make use of the notation introduced in Section 2.1.3. We take the perspective of one fixed receiver $k' \in \mathcal{K}$ (pictured on the left) for which we compute the SINR on all spatial layers $m = 1 \dots M_{k'}$ it receives from its desired transmitter $t' = t(k') \in \mathcal{T}$. In the following, we will just write t' and use the $t(k)$ mapping only when we refer to the transmitter of some other user $k \neq k'$. For the sake of simplicity, we omit the explicit time and frequency dependency of, for example, the channel matrices so that in this context the variable t always refers to the transmitter index. In a typical system-level simulation, the SINR computations discussed below would be repeated for each time and frequency unit (e.g., a TTI and a PRB) using the corresponding channels and resource allocation decisions. The desired transmitter t' transmits $M_{k'}$ spatial layers with complex baseband modulation symbols $s_{k',1}, \dots, s_{k',M_{k'}}$ by applying a precoding matrix $P_{k'} \in \mathbb{C}^{N_{\text{Tx},t'} \times M_{k'}}$ over an $N_{\text{Rx},k'} \times N_{\text{Tx},t'}$ MIMO channel $H_{t'k'}$. The receiver employs a receive filter $W_{k'}^H \in \mathbb{C}^{M_{k'} \times N_{\text{Rx},k'}}$ to recover the spatial layers as $\tilde{s}_1, \dots, \tilde{s}_{M_{k'}}$. In Fig. 4.2 transmitter t' transmits two spatial layers over a 2×2 channel by applying a suitable precoding matrix $P_{k'} \in \mathbb{C}^{2 \times 2}$. Another transmitter $t = t(k)$ transmits, in the example, just a single spatial layer precoded by a precoding vector $P_k \in \mathbb{C}^{2 \times 1}$ causing undesired interference to the receiver k' .

4.2.1. SINR Computation Per MIMO Layer

In a multi-cell scenario, not only the desired signal layers $\mathbf{s}_{k'} \in \mathbb{C}^{M_{k'}}$ from the desired transmitter t' , but also a number of interfering signals from other transmitters $t \in \mathcal{T} \setminus \{t'\}$, or even interfering intra-cell transmissions from the desired transmitter destined for other users $k \in \mathcal{D}_{t'}$ arrive at a receiver. Figure 2.3 on page 23 illustrates the scenario in the downlink and uplink of a cellular system. The following discussion of how SINRs per MIMO layer are modeled, mostly applies to the LTE downlink (because uplink MIMO was only recently introduced as a feature in Rel-10), but it is also valid for the uplink.

The antennas of receiver k' receive the superposition $\mathbf{y}_{k'} \in \mathbb{C}^{N_{\text{Rx},k'}}$ of all transmitted signals and noise $\mathbf{n}_{k'} \in \mathbb{C}^{N_{\text{Rx},k'}}$ with $[\mathbf{n}_{k'}]_i \sim \mathcal{CN}(0, \sigma_{\text{th},k'}^2)$ as already shown in (2.25) on page 22. A linear receiver filters the received vector $\mathbf{y}_{k'}$ with some filter matrix $\mathbf{W}_{k'}^{\text{H}}$ to obtain the filtered vector $\tilde{\mathbf{s}}_{k'}$, which it then maps component-wise to the estimated transmit vector $\hat{\mathbf{s}}_{k'}$ as described in Section 2.2.2. Equation (4.2) shows that the filtered vector $\tilde{\mathbf{s}}_{k'}$ also contains filtered contributions from the intra- and inter-cell interference as well as from the thermal noise:

$$\begin{aligned} \tilde{\mathbf{s}}_{k'} &= \mathbf{W}_{k'}^{\text{H}} \mathbf{y}_{k'} \\ &= \underbrace{\sqrt{P_{\text{tx},k'}} \mathbf{W}_{k'}^{\text{H}} \mathbf{H}_{t'k'} \mathbf{P}_{k'} \mathbf{s}_{k'}}_{\text{filtered desired signal}} + \underbrace{\sum_{k \in \mathcal{D}_{t'} \setminus \{k'\}} \sqrt{P_{\text{tx},k}} \mathbf{W}_{k'}^{\text{H}} \mathbf{H}_{t'k'} \mathbf{P}_k \mathbf{s}_k}_{\text{filtered intra-cell interference}} + \\ &\quad \underbrace{\sum_{t \in \mathcal{T} \setminus \{t'\}} \sum_{k \in \mathcal{D}_t} \sqrt{P_{\text{tx},k}} \mathbf{W}_{k'}^{\text{H}} \mathbf{H}_{tk'} \mathbf{P}_k \mathbf{s}_k}_{\text{filtered inter-cell interference}} + \underbrace{\mathbf{W}_{k'}^{\text{H}} \mathbf{n}_{k'}}_{\text{filtered noise}} \end{aligned} \quad (4.2)$$

If we sort the contributions to $\tilde{\mathbf{s}}_{k'}$ as in (4.2) above, we can derive the following expression for the SINR of each spatial layer m received from transmitter t' by user k' where $\mathbf{R}_{\text{IPN},k'}$ denotes the interference plus noise covariance matrix, which we introduced in (2.26) on page 23:

$$\text{SINR}_{k',m} = \frac{\overbrace{P_{\text{tx},k'} \left| \left[\mathbf{W}_{k'}^{\text{H}} \mathbf{H}_{t'k'} \mathbf{P}_{k'} \right]_{m,m} \right|^2}^{\text{signal power}}}{\underbrace{\sum_{\substack{i=1 \\ i \neq m}}^{M_{k'}} P_{\text{tx},k'} \left| \left[\mathbf{W}_{k'}^{\text{H}} \mathbf{H}_{t'k'} \mathbf{P}_{k'} \right]_{m,i} \right|^2}_{\text{inter-layer interference}} + \underbrace{\left[\mathbf{W}_{k'}^{\text{H}} \mathbf{R}_{\text{IPN},k'} \mathbf{W}_{k'} \right]_{m,m}}_{\text{filtered intra- and inter-cell interference}}}. \quad (4.3)$$

We can understand (4.3) by considering the $M_{k'} \times M_k$ matrix $\mathbf{W}_{k'}^{\text{H}} \mathbf{H}_{tk'} \mathbf{P}_k$ which maps the spatial layers in \mathbf{s}_k (be it desired or undesired) to the layers $\tilde{\mathbf{s}}_{k'}$ that we want to receive. The m -th component of the desired signal $\mathbf{s}_{k'}$ is mapped to $\tilde{\mathbf{s}}_{k'}$ by multiplication

with the m -th column of $\mathbf{W}_{k'}^H \mathbf{H}_{t'k'} \mathbf{P}_{k'}$. Ideally, the symbol $[s_{k'}]_m$ of layer m in $s_{k'}$ should only contribute to the m -th filtered signal layer $[\tilde{s}_{k'}]_m$. Hence, the effective gain experienced by the m -th desired layer $[s_{k'}]_m$ is $|\left[\mathbf{W}_{k'}^H \mathbf{H}_{t'k'} \mathbf{P}_{k'}\right]_{m,m}|^2$, which forms the nominator of (4.3). In practice, no perfect diagonalization is achieved so that the symbol of the m -th desired layer $[s_{k'}]_m$ causes inter-layer interference to the other spatial layers while it also suffers from inter-layer interference caused by the other layers. In fact, all entries in row m and column $i \neq m$ of $\mathbf{W}_{k'}^H \mathbf{H}_{t'k'} \mathbf{P}_{k'}$ cause inter-layer interference to the m -th layer and are thus to be found in the denominator. The second term in the denominator accounts for the effect of filtered intra- and inter-cell interference and noise as represented by the IPN covariance matrix $\mathbf{R}_{\text{IPN},k'}$ for user k' . The $[\mathbf{W}_{k'}^H \mathbf{R}_{\text{IPN},k'} \mathbf{W}_{k'}]_{m,m}$ expression allows for a concise notation, but the interference contribution can also be understood in the same way as described for the desired transmission. Let us consider a single interfering transmission from another transmitter $t \neq t'$ to some user k in a neighboring cell. We can then write the interference contribution $I_{k',m}^{(t \rightarrow k)}$ to the desired layer m at user k' as the sum of the contributions from all interfering layers $i = 1 \dots M_k$ that t transmits to user k . That sum for the single considered interfering transmission can be written using the receive covariance (in parentheses) of the interfering transmission from t to k as seen by user k' :

$$I_{k',m}^{(t \rightarrow k)} = \sum_{i=1}^{M_k} P_{\text{tx},k} \left| \left[\mathbf{W}_{k'}^H \mathbf{H}_{tk'} \mathbf{P}_k \right]_{m,i} \right|^2 = \left[\mathbf{W}_{k'}^H \left(P_{\text{tx},k} \mathbf{H}_{tk'} \mathbf{P}_k \mathbf{P}_k^H \mathbf{H}_{tk'}^H \right) \mathbf{W}_{k'} \right]_{m,m}. \quad (4.4)$$

The IPN covariance matrix $\mathbf{R}_{\text{IPN},k'}$ as defined in (2.26) contains all the intra- and inter-cell sources of interference and noise so that we obtain the concise representation as shown in (4.3).

4.2.2. Linear MIMO Receive Filters

The SINR model in (4.3) allows for arbitrary precoding matrices \mathbf{P} and filter matrices \mathbf{W} . In LTE, the precoders could be taken from the codebooks defined in the standard [36.211]. To adapt the receive filter \mathbf{W} to the current conditions, the receivers have to estimate the serving channel $\hat{\mathbf{H}}_{t'k'}$ and possibly also the covariance $\hat{\mathbf{R}}_{\text{IPN},k'}$ of the interference and noise. In Section 4.3 we will present ways to model the estimation errors when performing system-level simulations.

In the LTE specifications the spatial processing used at the receiver is not explicitly standardized so that vendors can differentiate themselves from the competition by implementing more sophisticated algorithms. Thus, for simulation-based system performance evaluations, assumptions regarding the receiver capabilities have to be made. Two baseline linear receivers (see also Section 2.2.2) are commonly used [36.829] for LTE Rel-8. The simpler one is the MRC receiver that combines the signals from multiple antennas by aligning the filter $\mathbf{W}_{k'}^{(\text{MRC})}$ to the estimated effective serving channel $\hat{\mathbf{H}}_{t'k'} \mathbf{P}_{k'}$:

$$\mathbf{W}_{k'}^{(\text{MRC})} = \hat{\mathbf{H}}_{t'k'} \mathbf{P}_{k'}. \quad (4.5)$$

The MRC receiver is only used for a single layer (rank-1) transmission. The other LTE Rel-8 MS baseline receiver is the linear MMSE or LMMSE receiver that allows to suppress the inter-layer, intra-cell, and inter-cell interference depending on what level of interference covariance knowledge is available at the receiver. It can be used for spatial multiplexing transmission modes where multiple layers are received:

$$\mathbf{W}_{k'}^{(\text{MMSE})} = \left(P_{\text{tx},k'} \hat{\mathbf{H}}_{t'k'} \mathbf{P}_{k'} \mathbf{P}_{k'}^H \hat{\mathbf{H}}_{t'k'}^H + \hat{\mathbf{R}}_{\text{IPN},k'} \right)^{-1} \hat{\mathbf{H}}_{t'k'} \mathbf{P}_{k'}. \quad (4.6)$$

The assumption for LTE Rel-8 receivers is that they are able to estimate the serving channel $\hat{\mathbf{H}}_{t'k'}$ and hence the covariance $\hat{\mathbf{H}}_{t'k'} \mathbf{P}_{k'} \mathbf{P}_{k'}^H \hat{\mathbf{H}}_{t'k'}^H$ of the effective serving channel but have only very limited knowledge of the IPN covariance $\hat{\mathbf{R}}_{\text{IPN},k'}$. In [R1-110586] the 3GPP defines a baseline MMSE receiver (called MMSE option 1) that is only aware of a diagonal IPN covariance matrix of all incoming signals that are not its desired signal:

$$\hat{\mathbf{R}}_{\text{IPN},k'} = \text{diag}(\sigma_{k',r}^2). \quad (4.7)$$

Here, $\sigma_{k',r}^2$ is the total noise and interference power per receive antenna r . That means, the receiver is able to detect different interference and noise power levels per antenna, but otherwise is not aware of the spatial interference structure.

In the unrealistic case that *perfect* IPN covariance knowledge is available, that is, if $\hat{\mathbf{R}}_{\text{IPN},k'} = \mathbf{R}_{\text{IPN},k'}$ as defined in (2.26), it can be shown (see, e.g., [MF11]) that the MMSE receiver in (4.6) becomes the *interference rejection combining* (IRC) or *optimum combining receiver* as introduced by Winters [Win84]:

$$\mathbf{W}_{k'}^{(\text{IRC})} = \left(P_{\text{tx},k'} \hat{\mathbf{H}}_{t'k'} \mathbf{P}_{k'} \mathbf{P}_{k'}^H \hat{\mathbf{H}}_{t'k'}^H + \mathbf{R}_{\text{IPN},k'} \right)^{-1} \hat{\mathbf{H}}_{t'k'} \mathbf{P}_{k'}. \quad (4.8)$$

The performance gains obtained by using different spatial processing techniques at the receiver will be evaluated and presented in Chapter 5. There, we will refer to the MMSE filter $\mathbf{W}_{k'}^{(\text{MMSE})}$ based on a diagonal IPN covariance matrix as the *MMSE-MRC* filter. The MMSE filter $\mathbf{W}_{k'}^{(\text{IRC})}$ based on a perfect IPN covariance matrix will be called the *IRC receiver* in single-layer contexts, whereas in spatial multiplexing settings, we will refer to it as the *MMSE-IRC* filter.

4.2.3. Transmit Diversity using Space-Frequency Block Codes

In Chapter 2 we introduced space-frequency block code (SFBC) transmit diversity, which is used in LTE transmission mode 2 and as the fallback transmission scheme for any transmission mode. It is also used for the transmission of the PDCCH if more than one antenna port is available. Due to its central role in LTE, we will shortly introduce the theory behind this transmission scheme and show how it can be cast to fit into the linear MIMO receiver framework presented earlier in this chapter.

The Alamouti Transmit Diversity Scheme

In contrast to the multi-antenna transmission techniques studied above, the aim of SFBC is not to enable the transmission of multiple independent spatial layers (*multiplexing gain*), or to achieve an antenna array gain (*beamforming gain*), but to make the transmission more robust against fading, that is, to obtain a *diversity gain*. The scheme was originally introduced by Alamouti in his famous 1998 paper [Ala98] as a *space-time block code* (STBC) for two Tx antennas in the time-domain. But Alamouti already mentioned the possibility to apply it in the frequency domain as implemented in LTE. The previously presented linear precoding schemes map a transmit symbol s_i on each subcarrier f_{idx} to multiple antennas. With SFBC, instead, consecutive symbols s_1 and s_2 are transmitted independently on adjacent subcarriers (SCs) f_1 and f_2 from Tx antennas 1 and 2 as shown in (4.9) [36.211, Sec. 6.3.4.3] where $(\cdot)^*$ denotes complex conjugation:

$$\begin{array}{c} \text{Tx Antenna 1} \\ \text{Tx Antenna 2} \end{array} \begin{array}{cc} \text{odd SC} & \text{even SC} \\ \left(\begin{array}{cc} x_1(f_1) & x_1(f_2) \\ x_2(f_1) & x_2(f_2) \end{array} \right) & = \frac{1}{\sqrt{2}} \left(\begin{array}{cc} s_1 & s_2 \\ -s_2^* & s_1^* \end{array} \right) \end{array} \quad (4.9)$$

With a transmission model as shown in (2.11) and in Fig. 2.2, we can write the received signals $y_r(f_{\text{idx}})$ at receive antenna r and subcarrier f_{idx} as

$$\text{odd SC: } y_r(f_1) = \frac{1}{\sqrt{2}} (h_{r,1}(f_1)s_1 - h_{r,2}(f_1)s_2^*) + n_r(f_1) \quad \text{and} \quad (4.10)$$

$$\text{even SC: } y_r(f_2) = \frac{1}{\sqrt{2}} (h_{r,1}(f_2)s_2 + h_{r,2}(f_2)s_1^*) + n_r(f_2), \quad (4.11)$$

where $h_{r,s}(f_{\text{idx}}) = [\mathbf{H}_{r's'}(f_{\text{idx}})]_{r,s}$ is the channel transfer function between Tx antenna s and Rx antenna r on subcarrier f_{idx} and $n_r(f_{\text{idx}})$ is the noise received on subcarrier f_{idx} . By conjugating the signal received on subcarrier f_2 in (4.11):

$$y_r^*(f_2) = \frac{1}{\sqrt{2}} (h_{r,1}^*(f_2)s_2^* + h_{r,2}^*(f_2)s_1) + n_r^*(f_2), \quad (4.12)$$

which has no impact on our ability to detect the transmitted symbol, we can cast equations (4.10) and (4.12) as $\bar{\mathbf{y}} = \bar{\mathbf{H}} \frac{1}{\sqrt{2}} (s_1 \ s_2^*)^T + \bar{\mathbf{n}}$ in a way that is analogous to (2.11). To facilitate this, we introduce the equivalent virtual received signal $\bar{\mathbf{y}}$, noise vector $\bar{\mathbf{n}}$, and channel matrix $\bar{\mathbf{H}}$ as follows [B B05]:

$$\underbrace{\begin{pmatrix} y_r(f_1) \\ y_r^*(f_2) \end{pmatrix}}_{\bar{\mathbf{y}}} = \underbrace{\begin{pmatrix} h_{r,1}(f_1) & -h_{r,2}(f_1) \\ h_{r,2}^*(f_2) & h_{r,1}^*(f_2) \end{pmatrix}}_{\bar{\mathbf{H}}} \underbrace{\frac{1}{\sqrt{2}} \begin{pmatrix} s_1 \\ s_2^* \end{pmatrix}}_{\bar{\mathbf{s}}} + \underbrace{\begin{pmatrix} n_r(f_1) \\ n_r^*(f_2) \end{pmatrix}}_{\bar{\mathbf{n}}}. \quad (4.13)$$

This notation reveals the beauty of the Alamouti precoding scheme: if we assume that the channel is constant over two adjacent subcarriers, that is $h_{r,1} := h_{r,1}(f_1) = h_{r,1}(f_2)$ and $h_{r,2} := h_{r,2}(f_1) = h_{r,2}(f_2)$, then the equivalent channel matrix $\bar{\mathbf{H}}$ is a scaled unitary

matrix for which $\bar{\mathbf{H}}^H \bar{\mathbf{H}} = \bar{\mathbf{H}} \bar{\mathbf{H}}^H = \text{const} \cdot \mathbf{I}$ holds. We can then use the Hermitian transpose of $\bar{\mathbf{H}}$ to filter the received signal $\bar{\mathbf{y}}$ and obtain the processed symbols \tilde{s}_1 and \tilde{s}_2^* for the transmitted symbols without any interference between the symbols as

$$\begin{aligned}
 \begin{pmatrix} \tilde{s}_1 \\ \tilde{s}_2^* \end{pmatrix} &= \begin{pmatrix} h_{r,1} & -h_{r,2} \\ h_{r,2}^* & h_{r,1} \end{pmatrix}^H \begin{pmatrix} y_r(f_1) \\ y_r^*(f_2) \end{pmatrix} \\
 &= \frac{1}{\sqrt{2}} \begin{pmatrix} h_{r,1}^* & h_{r,2} \\ -h_{r,2}^* & h_{r,1} \end{pmatrix} \begin{pmatrix} h_{r,1} & -h_{r,2} \\ h_{r,2}^* & h_{r,1} \end{pmatrix} \begin{pmatrix} s_1 \\ s_2^* \end{pmatrix} + \begin{pmatrix} h_{r,1}^* & h_{r,2} \\ -h_{r,2}^* & h_{r,1} \end{pmatrix} \begin{pmatrix} n_r(f_1) \\ n_r^*(f_2) \end{pmatrix} \\
 &= \frac{1}{\sqrt{2}} \begin{pmatrix} h_{r,1}^* h_{r,1} + h_{r,2} h_{r,2}^* & -h_{r,1}^* h_{r,2} + h_{r,2} h_{r,1}^* \\ -h_{r,2}^* h_{r,1} + h_{r,1} h_{r,2}^* & h_{r,2}^* h_{r,2} + h_{r,1} h_{r,1}^* \end{pmatrix} \begin{pmatrix} s_1 \\ s_2^* \end{pmatrix} + \begin{pmatrix} h_{r,1}^* & h_{r,2} \\ -h_{r,2}^* & h_{r,1} \end{pmatrix} \begin{pmatrix} n_r(f_1) \\ n_r^*(f_2) \end{pmatrix} \\
 &= \frac{1}{\sqrt{2}} \begin{pmatrix} |h_{r,1}|^2 + |h_{r,2}|^2 & 0 \\ 0 & |h_{r,2}|^2 + |h_{r,1}|^2 \end{pmatrix} \begin{pmatrix} s_1 \\ s_2^* \end{pmatrix} + \begin{pmatrix} h_{r,1}^* & h_{r,2} \\ -h_{r,2}^* & h_{r,1} \end{pmatrix} \begin{pmatrix} n_r(f_1) \\ n_r^*(f_2) \end{pmatrix}.
 \end{aligned} \tag{4.14}$$

We see that, under the assumption that the channels are constant over the two adjacent subcarriers, the Alamouti scheme is able to completely suppress the interference between the two symbols. The amplitude of both symbols is scaled by an identical factor of $1/\sqrt{2} (|h_{r,1}|^2 + |h_{r,2}|^2)$ so that the power is scaled by $1/2 (|h_{r,1}|^2 + |h_{r,2}|^2)^2$. Note, however, that the noise is also amplified by the $\bar{\mathbf{H}}$ filter with an expected power scaling of $|h_{r,1}|^2 + |h_{r,2}|^2$ so that in relation (i.e., for the SNR) the expected power gain per symbol results in $1/2 (|h_{r,1}|^2 + |h_{r,2}|^2)$. In total, there is thus no transmit-side power gain but a diversity gain from ‘‘averaging’’ the channels, which prevents deep fades as long as they are not correlated between the two antenna pairs. SFBC only requires channel knowledge at the receiver but not at the transmitter. Thus, it is especially well-suited for situations where the transmitter has no or only outdated or unreliable channel knowledge, for example in the case of a fast-moving receiver. This use case is also the motivation to use SFBC instead of STBC in LTE. High user speeds result in a short channel coherence time and are detrimental to STBC but do not affect the channel coherence in the frequency domain.

For 2 Tx antennas the Alamouti scheme achieves the full rate of 2 symbols per 2 subcarriers, which are transmitted in an orthogonal way on 2 Tx antennas. For more than 2 antennas, it has been shown [TJC99] that with complex-valued modulation signals an extension of the orthogonal design at full rate and diversity order is not possible. In LTE, for 4 Tx antennas the Alamouti scheme is thus extended over 4 subcarriers with only partly active antennas as shown in (4.16) [36.211, Sec. 6.3.4.3]:

$$\begin{pmatrix} x_1(f_1) & x_1(f_2) & x_1(f_3) & x_1(f_4) \\ x_2(f_1) & x_2(f_2) & x_2(f_3) & x_2(f_4) \\ x_3(f_1) & x_3(f_2) & x_3(f_3) & x_3(f_4) \\ x_4(f_1) & x_4(f_2) & x_4(f_3) & x_4(f_4) \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} s_1 & s_2 & 0 & 0 \\ 0 & 0 & s_3 & s_4 \\ -s_2^* & s_1^* & 0 & 0 \\ 0 & 0 & -s_4^* & s_3^* \end{pmatrix}. \tag{4.16}$$

Modeling for SFBC in the Linear Receiver Framework

We can accommodate SFBC in the linear receiver framework of our link-to-system model by employing the equivalent notation of (4.13) extended to an arbitrary number of N_{Rx} receive antennas. For the 2 Tx antenna SFBC scheme, we get the following equivalent system that has $2 \times N_{\text{Rx}}$ virtual Rx antennas and combines 2 subcarriers:

$$\underbrace{\begin{pmatrix} y_1(f_1) \\ \vdots \\ y_{N_{\text{Rx}}}(f_1) \\ y_1^*(f_2) \\ \vdots \\ y_{N_{\text{Rx}}}^*(f_2) \end{pmatrix}}_{\bar{\mathbf{y}}} = \underbrace{\begin{pmatrix} h_{1,1}(f_1) & -h_{1,2}(f_1) \\ \vdots & \vdots \\ h_{N_{\text{Rx}},1}(f_1) & -h_{N_{\text{Rx}},2}(f_1) \\ h_{1,2}^*(f_2) & h_{1,1}^*(f_2) \\ \vdots & \vdots \\ h_{N_{\text{Rx}},2}^*(f_2) & h_{N_{\text{Rx}},1}^*(f_2) \end{pmatrix}}_{\bar{\mathbf{H}}} \underbrace{\begin{pmatrix} \frac{1}{\sqrt{2}} & 0 \\ 0 & \frac{1}{\sqrt{2}} \end{pmatrix}}_{\bar{\mathbf{P}}} \underbrace{\begin{pmatrix} s_1 \\ s_2^* \end{pmatrix}}_{\bar{\mathbf{s}}} + \underbrace{\begin{pmatrix} n_1(f_1) \\ \vdots \\ n_{N_{\text{Rx}}}(f_1) \\ n_1^*(f_2) \\ \vdots \\ n_{N_{\text{Rx}}}^*(f_2) \end{pmatrix}}_{\bar{\mathbf{n}}}. \quad (4.17)$$

With the $\bar{\mathbf{y}} = \bar{\mathbf{H}}\bar{\mathbf{P}}\bar{\mathbf{s}} + \bar{\mathbf{n}}$ equivalent model in (4.17), we can use our previous linear receiver models and software implementation and implement the Alamouti combining scheme [Ala98] by choosing the MRC filter, that is, the Hermitian transpose of the equivalent channel as shown in equations (4.5) and (4.14) as the receive filter. For the SINR model in (4.3), we need to provide an equivalent IPN covariance matrix $\bar{\mathbf{R}}_{\text{IPN}}$:

$$\bar{\mathbf{R}}_{\text{IPN}} := \begin{pmatrix} \mathbf{R}_{\text{IPN}}(f_1) & \mathbf{0} \\ \mathbf{0} & \mathbf{R}_{\text{IPN}}^*(f_2) \end{pmatrix}, \quad (4.18)$$

where $\mathbf{R}_{\text{IPN}}(f_1)$ and $\mathbf{R}_{\text{IPN}}(f_2)$ are the (conjugated) inter-cell interference covariance matrices of subcarriers f_1 and f_2 , respectively. As the symbols on different subcarriers can be assumed to be independent, the IPN covariance matrix (i.e., the IPN part of $\bar{\mathbf{y}}\bar{\mathbf{y}}^{\text{H}}$) is block-diagonal with off-diagonal blocks of zeros. Note that no special treatment is necessary to model the performance impact an SFBC transmission has on other concurrent transmissions. The interference generated by an SFBC transmission appears like two independent transmissions or a dual-layer transmission with a $1/\sqrt{2}$ scaled identity matrix precoder $\bar{\mathbf{P}}$ because on a per-subcarrier basis, the symbols within the pairs are independent.

To accommodate the LTE 4 Tx antenna SFBC scheme shown in (4.16), the equivalent virtual model given in equations (4.17) and (4.18) can be extended to a size of $4 \times N_{\text{Rx}}$ in a straightforward way as implemented in IMTaphy.

4.3. Modeling Channel Estimation Errors

The receiver in an LTE system has to estimate both its serving channel $\hat{\mathbf{H}}_{t'k'}$ and—if it performs some kind of spatial processing to suppress interference—also the IPN covariance matrix estimate $\hat{\mathbf{R}}_{\text{IPN},k'}$. In an actual hardware implementation the estimation

process will never yield the exact channel or IPN covariance matrix so that the estimates $\hat{\mathbf{H}}_{\nu'k'}$ and $\hat{\mathbf{R}}_{\text{IPN},k'}$ will differ from the unknown actual values $\mathbf{H}_{\nu'k'}$ and $\mathbf{R}_{\text{IPN},k'}$. In a simulator, however, the channel matrices and thus also the covariance matrices are perfectly known. Hence, some extra effort has to be taken if the estimation errors occurring in an actual system are to be modeled. In the following, we will discuss different models that are implemented in `IMTAphy`. System-level simulations, however, are for comparison reasons often performed without channel estimation error models, especially for calibration purposes.

4.3.1. Modeling Serving Channel Matrix Estimation Errors

For the estimation of the serving channel matrix in the downlink of an LTE system, the MS makes use of the reference symbols transmitted by the BS from the different antenna ports. In Rel-8 there are only cell-specific reference symbols (see Fig. 2.12a on page 44), while later releases offer additional CSI-RS and DM-RS to support more antennas (see Fig. 2.12b). In the uplink, the channel estimation is done based on the demodulation reference symbols (when data is transmitted) or based on the sounding reference symbols, as shown on the right in Fig. 2.5.

Following an approach presented in [MF11, Sec. 10.2.3], we can model the channel estimation of the serving channel $\mathbf{H}_{\nu'k'}$ as tainted by an additive error matrix $\mathbf{E}_{k'}$ of identical dimensions and containing circularly symmetric complex Gaussian noise with zero mean:

$$\hat{\mathbf{H}}_{\nu'k'} = \mathbf{H}_{\nu'k'} + \mathbf{E}_{k'}. \quad (4.19)$$

In the LTE downlink the reference symbols used for channel estimation are protected from intra-cell interference because no data is transmitted in them and antenna ports are separated (see Fig. 2.5). Neighboring cells use different resource grid positions for their reference symbols so that the inter-cell interference to reference symbols mainly stems from REs carrying PDSCH data. As suggested in [MF11], we can thus model the random error matrix $\mathbf{E}_{k'}$ to have the same covariance as the general PDSCH IPN covariance but scaled by the inverse of a processing gain G_{est} that depends on the number of reference symbols available for estimation per Tx antenna and PRB and TTI:

$$E \left[\mathbf{E}_{k'} \mathbf{E}_{k'}^H \right] = \frac{1}{G_{\text{est}}} \mathbf{R}_{\text{IPN},k'}. \quad (4.20)$$

In our system-level simulator, we generate appropriate error matrices after performing a *Cholesky decomposition* $\mathbf{R}_{\text{IPN},k'} = \mathbf{L}\mathbf{L}^H$ of the perfect IPN covariance matrix $\mathbf{R}_{\text{IPN},k'}$ by multiplication with randomly generated circularly symmetric complex normal matrices $\mathbf{A} \sim \mathcal{CN}(\mathbf{0}, \mathbf{I})$:

$$\mathbf{E}_{k'} = \frac{1}{\sqrt{G_{\text{est}}}} \mathbf{L}\mathbf{A}, \quad (4.21)$$

so that with \mathbf{L} being a Cholesky factor and \mathbf{A} chosen to have an expected identity

matrix covariance matrix, we obtain the desired expected covariance matrix

$$E[\mathbf{E}_{k'}\mathbf{E}_{k'}^H] = E\left[\frac{1}{G_{\text{est}}}\mathbf{L}\mathbf{A}\mathbf{A}^H\mathbf{L}^H\right] = \frac{1}{G_{\text{est}}}\mathbf{R}_{\text{IPN},k'}. \quad (4.22)$$

According to the same authors [MF11, Sec. 10.2.3], the processing gain in an LTE system is in the range from 3 dB to 12 dB depending on how dispersive the channel is (i.e., how many CRS resource elements experience similar-enough channel conditions so that a processing gain can be realized given uncorrelated interference). In our implementation, we do not consider correlations of the random error terms between PRBs in the frequency domain or between TTIs in the time domain other than those due to similarities in the IPN covariance structure. Other channel estimation error models exist in the literature, for example, focusing on CoMP schemes [CDF+11].

4.3.2. Modeling IPN Covariance Matrix Estimation Errors

Now we turn to the estimation of the IPN covariance matrix, which is needed, for example, for the inter-cell interference-aware MMSE receivers introduced in (4.8). Here, we will restrict the discussion to the downlink, that is, we will only consider the estimation task for the receivers in the mobile stations. If the effective channels to all interferers could be well estimated, it would be possible to compute the IPN covariance matrix from the sum of the covariances of each interferer. However, that much knowledge about interfering inter-cell channels is hardly available in practice, and to suppress the inter-cell interference, it is also not necessary to distinguish between individual interferers. Instead, the covariance can be estimated directly from the received signal vectors $\mathbf{y}_{k'} \in \mathbb{C}^{N_{\text{Rx},k'}}$ containing the superposition of all signals received at the receiver's antenna array. Following the notation in [36.829] we express the estimation of a spatial covariance matrix as

$$\hat{\mathbf{R}}_{\text{IPN},k'} = \frac{1}{N_{\text{Sample}}} \sum_{n=1}^{N_{\text{Sample}}} \tilde{\mathbf{y}}_{k',n} \tilde{\mathbf{y}}_{k',n}^H, \quad (4.23)$$

where N_{Sample} is the number of samples and $\tilde{\mathbf{y}}_k$ is the received signal vector from which the contributions of the own cell have been canceled:

$$\tilde{\mathbf{y}}_{k'} = \mathbf{y}_{k'} - \hat{\mathbf{H}}_{l'k'} \mathbf{P}_{k'} \mathbf{s}_{k'}. \quad (4.24)$$

To be able to cancel the contribution from the own cell¹, the receiver already needs to have estimated the serving channel $\hat{\mathbf{H}}_{l'k'}$ and—most important—it already needs to

¹Note that if the MMSE filter is computed based on the received REs carrying the data to detect, the total covariance $P_{\text{tx},k'} \hat{\mathbf{H}}_{l'k'} \mathbf{P}_{k'} \mathbf{P}_{k'}^H \hat{\mathbf{H}}_{l'k'}^H + \hat{\mathbf{R}}_{\text{IPN},k'}$ needed in (4.8) could be estimated directly from the vector \mathbf{y}_k as defined in (2.25). However, here we want to derive only the inter-cell interference and noise part of the covariance of the received signal. The main motivation for separating channel and IPN covariance estimation is that for the feedback computation, the covariance of the desired signal is still unknown and subject to the precoder selection process.

know which symbol vector $\mathbf{s}_{k'}$ was transmitted. Unless some kind of iterative decision feedback technique is employed, this makes it impossible to use PDSCH data-carrying REs for estimating the IPN covariance in the same PRB and TTI. Thus, the IPN covariance estimation has to be performed on the a priori known CRS REs. Outside the control region, which, in a synchronized system, does not experience the same interference as the REs used for data transmission, there are 6–16 CRS REs available per PRB and TTI for this purpose, depending on the number of antenna ports used by the BS (see Fig. 2.12). Note that, unless appropriate system-wide coordination is performed, the inter-cell interference can vastly differ even between directly adjacent PRBs and TTIs due to different scheduling and precoding decisions in neighboring cells.

As described in [36.829; R1-111031] the sum over N_{Sample} samples from (4.23) follows a complex Wishart distribution with N_{Sample} degrees of freedom and perfect IPN covariance $\mathbf{R}_{\text{IPN},k'}$:

$$\sum_{n=1}^{N_{\text{Sample}}} \tilde{\mathbf{y}}_{k',n} \tilde{\mathbf{y}}_{k',n}^H \sim W_{N_{\text{Rx},k'}}(\mathbf{R}_{\text{IPN},k'}, N_{\text{Sample}}). \quad (4.25)$$

Consequentially, in TR [36.829, Sec. 8.2] the 3GPP proposed to model the covariance estimation errors for the system-level evaluation of advanced receivers in Rel-11 and beyond using the Wishart-distribution-based model. In IMT_{Aphy}, we adopt this model and generate the complex Wishart matrix $N_{\text{Sample}} \cdot \hat{\mathbf{R}}_{\text{IPN},k'}$ of size $N_{\text{Rx},k'} \times N_{\text{Rx},k'}$ using a so-called *Bartlett decomposition* (see [36.829; SH72] and [Bar33] cited in [SH72]) from a lower triangular matrix \mathbf{B} and the Cholesky factor \mathbf{L} :

$$\hat{\mathbf{R}}_{\text{IPN},k'} := \frac{1}{N_{\text{Sample}}} \mathbf{L} \mathbf{B} \mathbf{B}^H \mathbf{L}^H \quad \text{with } \mathbf{R}_{\text{IPN},k'} = \mathbf{L} \mathbf{L}^H. \quad (4.26)$$

The matrix \mathbf{B} is generated from random coefficients $c_i \sim \chi^2(2(N_{\text{Sample}} - i + 1))$, following a χ^2 -distribution with $2(N_{\text{Sample}} - i + 1)$ degrees of freedom, and circularly symmetric standard normal random coefficients $n_{i,j} \sim \mathcal{CN}(0, 1)$ [36.829; SH72]:

$$\mathbf{B} := \begin{pmatrix} \sqrt{c_1/2} & 0 & 0 & \cdots & 0 \\ n_{2,1} & \sqrt{c_2/2} & 0 & \cdots & 0 \\ n_{3,1} & n_{3,2} & \sqrt{c_3/2} & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ n_{N_{\text{Rx},k'},1} & n_{N_{\text{Rx},k'},2} & n_{N_{\text{Rx},k'},3} & \cdots & \sqrt{c_{N_{\text{Rx},k'}/2}} \end{pmatrix}. \quad (4.27)$$

The estimation accuracy depends on the number of samples. With increasing N_{Sample} , the estimated IPN covariance matrix $\hat{\mathbf{R}}_{\text{IPN},k'}$ converges to the perfect IPN covariance $\mathbf{R}_{\text{IPN},k'}$. In Chapter 5, we will discuss and present the performance of advanced receivers under different assumptions regarding the number of available covariance estimation samples.




Figure 4.3.: QPSK MMIB-based effective SINR computation example [EMD]

4.4. Link Performance Model

In the previous section, we introduced the link quality model that allows us to model the transmission conditions for an LTE transport block in a multi-cell system by providing a number of SINR values on the occupied physical resources as shown in Fig. 4.1. Now, we will discuss the link performance modeling part that derives a block error probability from these SINRs. As illustrated in Fig. 4.1, the first step is the computation of a single effective SINR from the individual SINRs, followed by a table lookup of the BLER based on pre-computed lookup tables.

4.4.1. SINR Compression to Compute Effective SINR Values

There are two main methods used in the literature to combine multiple SINRs that are encountered by a code block into a single effective SINR. Both the *exponential effective SINR mapping* (EESM) and the *mutual information effective SINR mapping* (MIESM) transform the individual SINR values using a suitable mapping $g(\cdot)$ to a domain where the transformed values can be arithmetically averaged. This average is then transformed back using the inverse mapping $g^{-1}(\cdot)$, see Brüninghaus et al. [BAS+05; EMD]. Schematically, for N_{SINR} individual values SINR_n , the effective SINR is computed as:

$$\text{SINR}^{(\text{eff})} = g^{-1} \left(\frac{1}{N_{\text{SINR}}} \sum_{n=1}^{N_{\text{SINR}}} g(\text{SINR}_n) \right). \quad (4.28)$$

Here, we will focus on an MIESM variant that maps SINRs to the mean mutual information per bit (MMIB) domain. The MMIB is the BICM rate [CTB98] in an AWGN channel divided by the number of coded bits per modulation symbol for which the 802.16m evaluation methodology document [EMD; SZS07] provides numerical approximations. Fig. 4.3 exemplifies how the MMIB-based effective SINR computation process works. In the example, two SINR values (0 dB and 20 dB) are mapped to the

corresponding MMIB values (0.486 and 1) indicated on the right figure axis. In that domain, the arithmetic mean $(0.486 + 1)/2 = 0.743$ is computed and then mapped back to an effective SINR value of 3.33 dB. Numerical approximations of the MMIB mapping function for different modulations are given in the EMD [EMD; SZS07]. The mapping and its inverse are often implemented by means of a lookup table. Note that the reverse mapping for MMIB values of 1 is not well-defined so that effective SINRs are constrained to a certain range.

The rationale for mutual information-based averaging approaches is easy to understand. Each individual SINR_n value represents the conditions that some of the modulation symbols in the code block encounter during the transmission. Looking at a single symbol, the mutual information conveyed by that symbol is constraint by its modulation order. Even for arbitrarily-high SINRs, a QPSK symbol can carry a maximum of 2 bits, a 16-QAM symbol a maximum of 4 bits, and a 64-QAM symbol a maximum of 6 bits. For QPSK modulation, the maximum MMIB value of 1 is already reached at about 10 dB. Thus, increasing the SINR value of the QPSK symbol beyond 10 dB does not contribute additional mutual information. If no effective SINR technique was used, and if, for example, the SINR values of 0 dB and 20 dB in the example above were to be averaged directly in the SINR domain, the 20 dB value would be overrepresented, leading to a too-high effective SINR value.

In the extreme case, there could be one SINR value for each RE used by the transport block. But to save the computational overhead of modeling the channel and the post-receiver SINRs with such a high granularity, most system-level simulators including `IMTAphy` generate just one SINR per 1 ms TTI. In the frequency domain, `IMTAphy` allows to generate channel realizations and SINRs per 15 kHz subcarrier or only one value for a 180 kHz PRB. The modeling of just one value per TTI and PRB is justified because the dimensions of the smallest schedulable resource unit in LTE (1 ms TTI \times 180 kHz PRB) are chosen to be within the channel coherence time and bandwidth of most use cases. Also, the interference situation is constant during one TTI and one PRB in a synchronized system. For one transport block, we thus model either one SINR per PRB or per subcarrier for all of the PRBs and layers spanned by the transport block.

For the link-to-system model in the SC-FDMA uplink in LTE a slightly different aggregation approach has to be used. Although uplink transmissions are scheduled onto PRBs as in the downlink, those PRBs have to be adjacent because uplink transmissions actually happen on a wider single carrier. In contrast to the OFDMA downlink, the need for more sophisticated frequency domain equalization arises in the uplink [DPS11]. Usually, a model for MMSE-based equalization in the frequency domain is used that yields a single SINR value per symbol duration. Conventional effective SINR computation approaches must be used in addition if multiple SINRs of individual symbol durations during the TTI or multiple layers are modeled. According to 3GPP assumptions [R1-051352; R1-050718], the post-equalization $\text{SINR}^{(\text{MMSE-FDE})}$ can be computed from the SINR_n values on the N_{SINR} subcarriers as:

$$\text{SINR}^{(\text{MMSE-FDE})} = \frac{\Gamma^2}{N_{\text{SINR}} \cdot \Gamma - \Gamma^2} \quad \text{with} \quad \Gamma = \sum_{n=1}^{N_{\text{SINR}}} \frac{\text{SINR}_n}{\text{SINR}_n + 1}. \quad (4.29)$$




Figure 4.4.: Block error ratios for different MCSs and block lengths based on LTE CQIs

4.4.2. Block Error Modeling

Once an effective SINR is computed for a code block, the expected BLER can be looked up in pre-computed BLER tables like the ones plotted in Fig. 4.4. Of course, the error probability depends on the code rate, that is, the ratio of actual information bits versus the number of coded bits. A smaller code rate implies more redundancy bits and thus allows successful decoding in the presence of more (coded) bit errors.

Figure 4.4 shows BLER curves that were generated by the author during a detailed link-level-simulation campaign with the MATLAB-based link-level simulator [MWI+09] developed at the TU Vienna. These and some additional curves are stored as lookup tables in *IMT_{phy}*. The link-level simulations were performed assuming a single BS–MS link over a non-fading SISO AWGN channel. During the simulations, the SINR was varied over the whole range and for each SINR value, one transport block per TTI was transmitted for 5000 TTIs in a row. After these 5000 trials, the empirical BLER was computed from the ratio of erroneously received transport blocks. The link-level simulator [MWI+09] was slightly modified to vary the size of the transport block from 50 to the maximum size of 6144 information bits in LTE (or as close to these numbers as possible given the minimum number of bits per MCS and PRB). The modulation scheme and code rates for the transmission of each transport block were chosen according to one of the 15 possible CQI values in LTE, see Table B.1 in the Appendix.

As can be seen in Fig. 4.4, different modulations and code rates lead to different threshold SINRs at which the BLERs start to fall from 1 towards 0. Different block lengths also influence the decoding chances because Turbo coding works better for larger code blocks. However, small block lengths already show BLERs < 1 at lower SINRs because chances are that sometimes the few modulation symbols per block encounter only small random AWGN noise realizations that allow the block to be decoded (higher variance with few samples). In LTE the goal of the link adaptation mechanism is to estimate the channel conditions (i.e., the SINR that a transport block will encounter) in order to select the highest modulation and coding scheme that can still be received successfully. Figure 4.5a shows the expected modulation and coding efficiency that

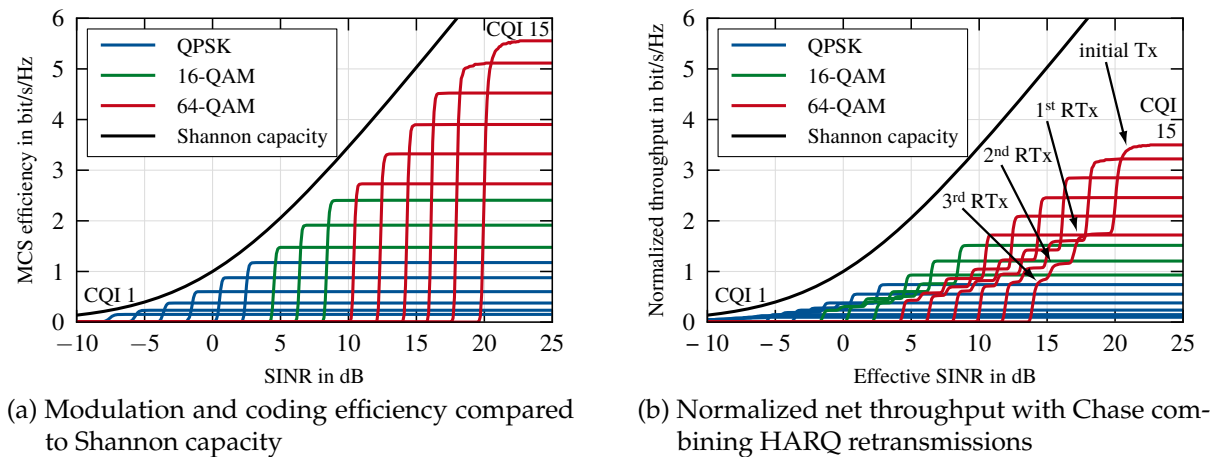


Figure 4.5.: Performance of different MCSs based on the 15 LTE CQIs

results from the BLER curves in Fig. 4.4 and the effective code rate for the 15 LTE CQIs given in Table B.1. We see that with a proper link adaptation, it is possible to follow the Shannon capacity by selecting the most suitable modulation and coding scheme. The gap of about 25%–30% between the envelope of the curves for the individual CQIs and the Shannon bound represents the losses e) to g) according to Table 2.2. As already outlined in Section 2.3.1, the modulation and coding efficiencies shown in Fig. 4.5a are not achievable in practice due to the various overheads. Taking a typical overhead situation for LTE transmission mode 1 (i.e., $L = 3$ control region symbols, CRS for one antenna port, cyclic prefix, and guard bands) into account, the normalized throughputs as depicted in Fig. 4.5b can be achieved in an LTE system. The figure also shows that with the help of HARQ retransmissions and chase combining, the SINR threshold of each MCS can be lowered by 3 dB, 4.8 dB, or 6 dB for one, two, or three retransmissions (RTx), respectively (see the three steps in the curves for each CQI in Fig. 4.5b). The availability of HARQ allows to tune the link adaptation process to aim at a residual BLER of about 10% because, as experience shows, at this point the gains from using better MCSs outweigh the losses incurred by a retransmission.

In an LTE system-level simulation, many different code block sizes between 50 and 6144 bits are possible. Also, arbitrary code rates can result from the rate matching process that depends on the number of resource elements within the allocated PRBs that are not used for other purposes. Therefore, the link-to-system interface in IMTaphy performs an interpolation if the actual block size or the code rate lie in between two of the pre-computed curves. To look up an arbitrary code rate, the curves for the next higher and lower code rate are shifted in the SINR direction. To look up an arbitrary block length, a linear interpolation between the BLERs of two available block lengths is performed. The LTE rate matcher achieves code rates below the Turbo coder *mother code rate* of $\frac{1}{3}$ by repeating an according number of bits. For low code rates, the rate- $\frac{1}{3}$ curves are thus used and the SINRs are scaled to reflect the power gain from repetition coding. Similarly, the effect of a Chase combining HARQ retransmission can be modeled by adding the SINRs of the initial transmission and of each retransmission on a

PRB-to-PRB basis before computing the effective SINR. Some authors [WPS+08] scale the SINRs of each retransmission by a combining efficiency factor to reflect the imperfections of Chase combining. Since the suggested [WPS+08] efficiency of 95% is close to one, we do not account for an efficiency loss in IMT_{Aphy} . Modeling an *incremental redundancy* (IR) retransmission would involve a mixture between an adjustment of the code rate and a fractional scaling of the SINR depending on how many bits are retransmitted and how many new bits are provided by the IR retransmission.

4.5. 3GPP IMT-A System-Level Simulator Calibration

Building on the channel and link-to-system model introduced in the preceding sections, we are now able to look at a complete LTE system-level simulation. As an example, we consider the simulation assumptions and the results of the simulator calibration campaign that the 3GPP conducted to prepare for the submission of LTE-Advanced as an IMT-Advanced technology in 2009. The calibration campaign does not reflect the performance that can be expected from an actual LTE or LTE-Advanced system because only a very basic setup was chosen. After calibrating their simulation tools, the 3GPP demonstrated LTE's IMT-Advanced target fulfillment with simulation results generated using more realistic simulation assumptions, as documented in TR [36.814].

Similar to the channel model calibration discussed in Section 3.4, the IMT-Advanced calibration setup is well-defined and does not leave much room for vendor-specific optimizations like sophisticated schedulers or receivers. What it does allow, however, is to capture the effects that the fast-fading channel and the link-to-system modeling assumptions have on the performance metrics. For an independent implementation like IMT_{Aphy} , it can thus serve as a benchmark that allows a comparison with the simulation tools used by 3GPP members.

The IMT-Advanced calibration simulation assumptions are described in the annex of TR [36.814]. The four mandatory IMT-Advanced scenarios InH, UMa, UMi, and RMa are considered with all the scenario parameters like user speeds etc. chosen as described in Section 3.1. Inter-cell interference is explicitly modeled by performing the simulation with the same level of detail in each cell. The most important LTE system parameters are summarized in Table 4.1. We will briefly discuss the simulation assumptions separately for downlink and uplink in the following. Finally, in Section 4.5.3 we will present the results of the calibration simulations.

4.5.1. Downlink Simulation Assumptions

The DL simulation setup assumes a SIMO system with a single BS transmit antenna and two receive antennas at the mobile stations. The 10 users per cell (on average) are associated to their serving cells taking into account a 1 dB handover margin. The mobiles are configured to provide wideband CQI reports to their serving base station every 5 ms. As the base stations are configured to use LTE transmission mode 1 and are

Table 4.1.: 3GPP assumptions for IMT-Advanced simulator calibration [36.814]

Parameter	Value
Common:	
Bandwidth	50 PRBs (10 MHz) in UL and DL each (FDD) for UMa, UMi, and RMa and 100 PRBs for InH
Receiver type	MRC receiver
Channel estimation	perfect channel estimation
Feeder loss at the BS	0 dB, in contrast to the calibration in Section 3.4.1
HARQ	Chase combining with up to three retransmissions
Antenna configuration	1 Tx and 2 Rx vertically polarized antennas, see Fig. 3.4 $\lambda/2$ separation at MS and 10λ separation at BS
Downlink:	
Control region overhead	$L = 3$ out of 14 symbols occupied each TTI
CQI feedback	5 ms periodicity, error-free, wideband estimate from TTI n can be used in TTI $n + 6$
CRS configuration	one CRS antenna port for transmission mode 1
Uplink:	
Frequency domain equalization	MMSE model assuming perfect channel estimates
Power control	open-loop only with $P_0 = -106$ dBm, $\alpha = 1$
SRS overhead	1 out of 14 symbols occupied every 10 ms
PUCCH overhead	4 PRBs deducted from the system bandwidth
Link adaptation	ideal estimate from TTI n can be used for TTI $n + 7$

equipped with one transmit antenna only, they cannot perform any precoding. Accordingly, no rank or precoding matrix indicators are signaled by the mobiles. The network employs a very simple round-robin scheduler that assigns an equal share of resources to all users. In a strictly sequential order, one user after the other is assigned all PRBs during one TTI. The users get different data rates though because the scheduler performs link adaptation based on the wideband CQI status reports, which are expected to be available 6 ms after they are measured. For the reception, the mobiles use their two antennas to perform maximum ratio combining based on perfectly known serving channels. A fixed control region size of three symbol durations is modeled. In addition, six REs are occupied by single-port cell-specific reference symbols in each PRB and TTI. Every 10-th TTI, the primary broadcast channel occupies some of the six center PRBs, which are then also not available for data transmission. The chosen settings correspond to the first column with $L = 3$ in Table 2.3.

4.5.2. Uplink Simulation Assumptions

The uplink simulations also use a SIMO model, this time with a single transmit antenna at the mobile and two antennas at the BS, which are again used for MRC combining with perfect channel estimates. Compared to the downlink, there are three major differences in the uplink: First, in LTE there is an open-loop power control in the uplink with

an optional closed-loop component in addition. For the calibration simulations, only open-loop is assumed and configured in a way that each MS chooses the Tx power so that it compensates the pathloss. The BS can thus expect to receive all users with similar Rx power levels. Second, the round-robin scheduler assumed for calibration works differently: In all cells, the $N_{\text{PRB}} = 50 - 4 = 46$ available PRBs are divided equally among all associated users in each TTI. As shown on the right-hand side in Fig. 2.5 on page 26, every user is served with always the same PRB allocation each TTI, regardless of the transmission being a new transmission or a HARQ retransmission. This special variant of round-robin scheduling leads to a quite unrealistic interference situation in the uplink. The resulting SINR and throughput distributions do not exhibit the typical uplink interference dynamics, which we discussed in Section 2.1.3. Third, the overhead situation in the uplink is different from the downlink as shown in the right half of Fig. 2.5: There is no control region and no pattern of cell-specific reference symbols like in the downlink. Instead, all subcarriers for 2 out of 14 symbol durations per TTI are used for transmitting demodulation reference symbols and every 10 TTIs, one symbol duration is assumed to be used for an uplink SRS in this setup. In addition, in total 4 out of 50 uplink PRBs at the lower and upper edge of the system bandwidth are not available for transmissions because they are reserved for the PUCCH.

4.5.3. Simulator Calibration Results

The first step (“step 1A”) of the 3GPP calibration campaign focusing on the large-scale fading was already presented in Section 3.4.1. As part of the second step (“step 1C”), we now look at the CDFs of the post receiver-combining SINR for the downlink and uplink as shown in Figs. 4.6a and 4.6b for the four mandatory scenarios. The 3GPP chose to compare these post-combining SINR distributions because they are well-defined and allow to reveal differences in the generation of fading channels and link quality modeling. The SINR is averaged per-user in the linear scale over all PRBs used in the transmissions to and from a user. For the IMT_{Aphy} results, each drop was simulated for one second of simulation time and a total of 25 drops was generated for the outdoor scenarios. For the InH scenario 350 drops were used. That many InH drops are necessary because only 20 users are placed on the scenario per drop. In addition, the area is so small that, due to the spatially-correlated shadowing, no reliable average can be derived from a single drop even if there were more users.

At first glance, the downlink SINR CDFs plotted in Fig. 4.6a look similar to the geometry ones shown in Fig. 3.8b on page 77. But in contrast to the previous results, two receive antennas with MRC-combining are now used with a fast-fading channel. The resulting MRC gain of up to 3 dB is visible, when carefully comparing the non-fading SISO geometry CDF in Fig. 3.8b to the CDF in Fig. 4.6a. For example, the curves top out at about 20 dB instead of 17 dB as before. Another effect of working with fast-fading serving and interfering links is visible when closely examining the improvement in the InH CDF between Figs. 3.8b and 4.6a. For example, at the 20 % level of the CDFs, a gain of more than the 3 dB attributable to MRC-combining can be seen. The reason for this effect is the unique characteristic of the InH scenario where only one interferer exists

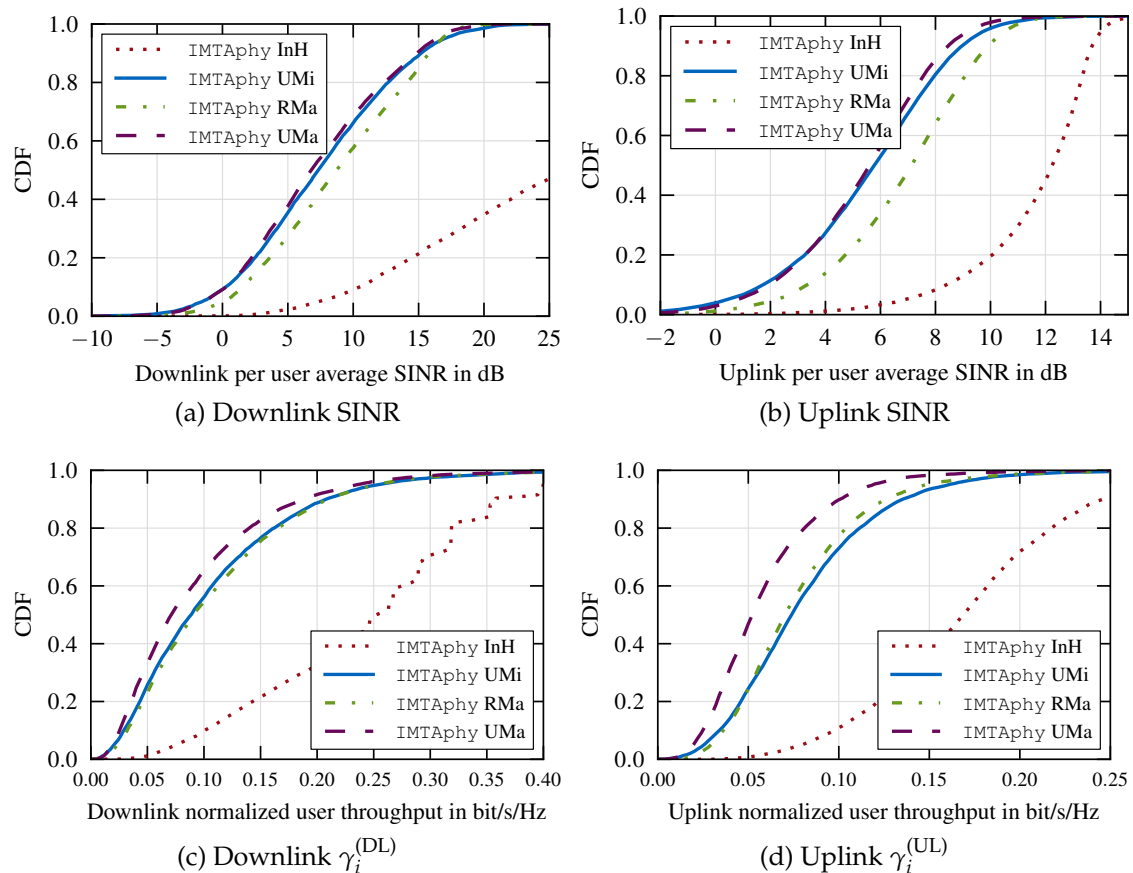


Figure 4.6.: Comparison of the CDFs of the average per-user post-receiver SINRs and normalized throughput for the IMT-A scenarios

whose interference is usually received much stronger than the thermal noise level. The noise is then irrelevant and the SINR (or SIR) is just the ratio of two χ^2 -distributed random variables: the fast-fading gains of the desired and of the interfering signal. In the InH scenario, the single interferer can exhibit deep fades leading to high SINR values that significantly increase the average. In the hexagonal scenarios with 56 interferers, it is highly unlikely that all interferers are faded at the same time so that the fast-fading of the signal power averages out over time and no such effect is visible.

The distribution of uplink SINRs shown in Fig. 4.6b is heavily influenced by the special power control settings and the round-robin variant used for calibration purposes. The visible SINR variance in the CDFs does not stem from pathloss effects because they are fully compensated by the power control mechanism. Neither is the fast-fading responsible for the SINR variance because the CDF shows the distribution of long-time average values. The main reason why the MSs achieve different average SINRs is the arbitrary distribution of favorable and unfavorable static interference constellations caused by the scheduler used for the calibration.

Figures 4.7a and 4.7b compare the DL and UL post-combining SINR distributions obtained with IMTAphy in the RMa scenario to the 3GPP average results as published in

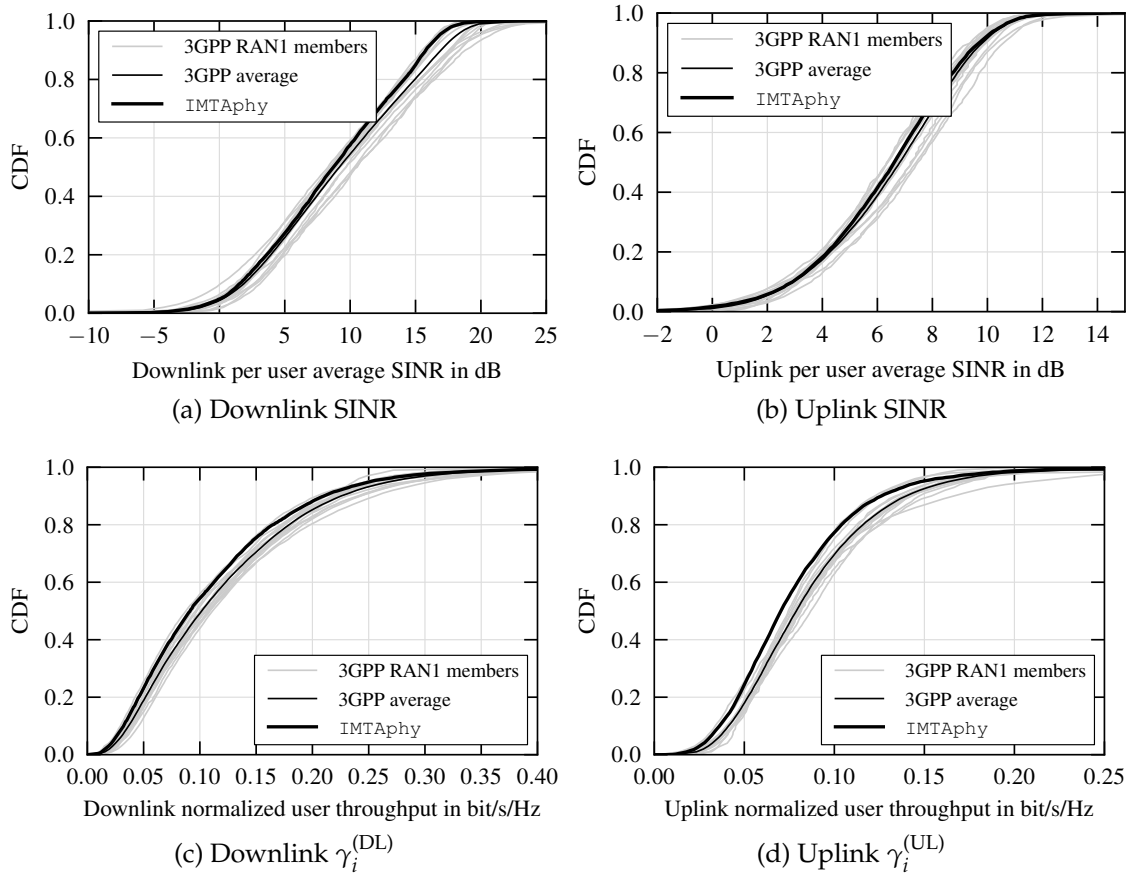


Figure 4.7.: CDFs of the avg. per-user post-receiver SINRs and normalized throughput in the RMa scenario compared to 3GPP RAN1 results [R1-092742]

TR [36.814]. In addition, the CDFs obtained by the individual contributors of the 3GPP RAN1 working group are also plotted based on results published in [R1-092742]. Compared to the non-fading SISO geometry SINR distribution in Fig. 3.8b on page 77, the results obtained from the fading channels with the respective link-to-system models exhibit a much wider spread with significant deviations from the average curve shown in TR [36.814]. One effect contributing to these deviations, which is especially visible in the RMa scenario shown in Fig. 4.7, seems to be that some companies model the fading to co-located base stations in the LoS case as uncorrelated. Most of the high geometry links exhibit LoS propagation (c.f. Table 3.3) with the two co-located BSs limiting the geometry to 17 dB due to the 20 dB forward-backward attenuation of the antenna pattern as discussed in Section 3.2.1. For LoS propagation, most power is then concentrated in the single LoS rays that arrive at identical angles from all co-located BSs. Thus, the channels are almost fully correlated, thereby preventing the MRC combining from realizing a 3 dB gain achievable with uncorrelated interference.

Finally, Figs. 4.6c and 4.6d show the per-user throughput CDFs for the four scenarios, and Figs. 4.7c and 4.7d compare the throughput in the RMa scenario achieved by IMTAphy to 3GPP results. Since the throughput performance directly depends on

the achievable SINRs, the picture is similar to the SINR case. The *IMTAphy* curve exhibits slightly lower throughput than most of the 3GPP RAN1 members, which can be explained by a more conservative link-to-system model. For example, some 3GPP members use IR HARQ, which can offer slightly better performance than the CC HARQ modeled by *IMTAphy*. Again, we observe a high variation in the reference results because besides the fading channel generation, all aspects of the link-to-system model as well as the RRM schemes contribute to the performance result.

We should note again that these calibration results are not representative for the performance of LTE. In Chapters 5 and 6 we will discuss how various RRM techniques can be employed to achieve higher downlink spectral efficiencies.

4.6. Summary and Contributions

In this last chapter on simulation model basics, we motivated the use of a so-called link-to-system model to facilitate large-scale system-level simulations with a feasible computational effort. After a review of potential alternatives, we presented the link-to-system model used in *IMTAphy*, which consists of a link quality and a link performance part. It maps the channel realizations generated by the channel model, as well as the link-adaptation and precoding decisions by the serving and interfering transmitters, to a transport block error probability at the intended receiver.

The link quality part of the model uses post-receiver SINRs to quantify the linear receiver performance. Supporting all LTE use cases, we discussed the per-layer SINR computation considering arbitrary linear receive filters, the use of linear precoders at the transmitter, and the interference covariance impact from precoded interfering transmissions. We also presented how to account for the Alamouti transmit diversity scheme used in LTE by relying on an equivalent virtual system model representation.

For their use in the performance part of the model, multiple SINR values (e.g., for different PRBs) provided by the link quality model have to be compressed into an effective SINR. This is done by mapping the downlink SINRs to and from the mutual information domain, where an arithmetic average can be computed. The link performance part then relies on pre-computed lookup tables that map the effective SINR value of a transport block and the applied modulation scheme and code rate into a block error value. We extended the conventionally used lookup curves by contributing a set of curves that also take the code block sizes into account.

As a further modeling aspect that sets *IMTAphy* and the evaluation presented in this thesis apart from other available tools, we presented ways to account for channel and interference covariance matrix estimation errors.

Finally, we discussed selected results of a rigorous calibration campaign that compares *IMTAphy*'s per-user post-receiver SINR and throughput distributions to those of 3GPP member companies. We could show that, in addition to the channel model, also the combination of channel model and link-to-system model deliver comparable results, which underlines the accuracy of our link-to-system model.

5. Performance of Fundamental RRM Schemes

So far, we have introduced the fundamentals of wireless communications in general and of LTE in particular (Chapter 2). Equipped with IMT-Advanced evaluation scenarios and channel models (Chapter 3), as well as link performance models (Chapter 4), in this chapter we can now evaluate the performance of LTE using `IMT_Aphy`. For the sake of brevity, the discussion will be focused on the downlink and only consider basic LTE Rel-8 RRM mechanisms. We will put a special emphasis on the interplay of different fundamental RRM methods and evaluate the dependencies on equipment configuration choices like antenna arrays and on external factors such as radio channel characteristics. In this chapter, we will focus on RRM schemes that primarily depend on and target at the conditions in the cell sector served by a base station. In the subsequent chapter, we will then look at mechanisms that take multiple cells and their inter-cell interference dependencies into account.

In the following, we will gradually introduce several RRM schemes as we increase the complexity of the considered system from a simple SISO link in Section 5.1 to a full-blown MIMO link in Section 5.5. At the same time, we will take a detailed look at transmission modes TM1, TM2, TM3, TM4, and TM6 and finally directly compare these five most important LTE transmission modes in Section 5.6. Throughout the discussion, we will highlight that it is not possible to quantify the performance (i.e., the spectral efficiency) of LTE with a single universal number: the spectral efficiency strongly depends on the employed RRM schemes, BS and MS configuration choices, the radio propagation characteristics of the deployment scenario, and the inter-cell interference situation.

5.1. Scheduling and Link Adaptation (SISO)

We begin our performance evaluation of RRM schemes by looking at a SISO link because with just a single antenna at the BS and MS, we can exclude any spatial processing from the discussion. Instead, we focus on the effects that basic resource and user scheduling, as well as link adaptation strategies (operating on quantized and delayed feedback), have on the spectral efficiency and throughput fairness in LTE.

Table 5.1 lists the most important simulation parameters that are common to the simulations discussed in the following. In addition to that, all parameters specific to the indicated IMT-Advanced deployment scenario (e.g., UMi, UMa, RMa, and InH) apply, see Tables 3.1 and 3.2 in Chapter 3. For each individual simulation result in the remainder of this chapter, we will only specify the parameters that are changed with respect to the general assumptions in Table 5.1.

Table 5.1.: General simulation parameters for the simulations in this chapter

Parameter	Value
Simulation Setup:	
Simulation time	1 s per drop, 250 ms thereof not evaluated (settling time)
Subcarriers	1 subcarrier per PRB modeled
Scenario Parameters:	
Bandwidth	10 MHz, i.e., 50 PRBs
BS antennas	1 purely vertically polarized antenna
MS antennas	1 purely vertically polarized antenna
Transmission mode	TM1
User speed	default for IMT-A scenario (e.g., 3 km/h for UMi)
Number of BS	57 with wraparound modeling for UMi, UMa, and RMa 2 base stations without wraparound for InH
Number of MS	on average $N_{MS} = 10$ users per cell
Feedback Configuration:	
Feedback delay	6 ms, error-free delivery
CQI update frequency	every 5 ms
CQI granularity	wideband (i.e., 50 PRBs per subband)
Scheduling and Link Adaptation:	
Link adaptation	fixed mapping, no offset
Schedulers	proportional fair and round-robin
HARQ	synchronous Chase combining HARQ
Overhead Assumptions:	
Control region overhead	$L = 3$ out of 14 symbols occupied each TTI
CRS ports	depending on number of BS antennas (see Fig. 2.12)

5.1.1. Channel-Adaptive Proportional Fair Downlink Scheduling

The basics of channel-adaptive user and resource scheduling and the necessary CSI feedback mechanisms in LTE have already been introduced in Section 2.2.3. Here, we will focus on some implementation details and, most importantly, on the performance under different circumstances.

Implementation Aspects of Round-Robin and Proportional Fair Schedulers

Figure 2.9 on page 39 illustrates the basic principles and differences between the above-mentioned scheduling variants. Round-robin is relatively straightforward to implement, but a few aspects have to be specified¹: unless otherwise stated, we implement the round-robin scheduling in the time-domain, that is, each user is allocated all PRBs

¹Note that we do not provide pseudocode and only describe the algorithms employed in IMT_{Aphy} verbally. For further details, the interested reader can refer to the complete IMT_{Aphy} source code, which is freely available on the Internet [IMT_{Aphy}LP].

and is not scheduled a new transmission again before all other active users have been scheduled as well. As indicated in Table 5.1, we perform HARQ retransmissions in a synchronous way. That means we retransmit once a NACK is received (i.e., 8 ms after the previous transmission attempt) without degrading the user in the round-robin queue. Compared to an implementation that only retransmits when a user is the next in line, our implementation can thus exhibit a certain bias towards lower SINRs and throughputs because users with retransmissions are scheduled more often.

Our proportional fair implementation follows the proportional fair metric introduced in (2.31) on page 39, but in (5.1) below we add a few more details. Instead of the instantaneous SINR in (2.31), which is not known to the BS in a real system, the BS has to rely on the CQI feedback. The CQI definition in LTE only provides 15 different CQI levels, each associated with an effective normalized data rate of $r(\text{CQI})$, see Table B.1 in the Appendix. The MS can be configured to provide frequency-selective feedback so that we have to distinguish between PRBs. In a TTI n and for transmission on PRB f_{idx} , the BS selects that user $k^*(f_{\text{idx}}, n)$, who maximizes the PF metric

$$k^*(f_{\text{idx}}, n) = \arg \max_{k \in \mathcal{K}} \frac{r(\text{CQI}(k, f_{\text{idx}}, n))}{\overline{\gamma}_{k,n}^\beta}, \quad (5.1)$$

where β is used to weight the influence of the average past throughput $\overline{\gamma}_{k,n}$. As we will see, the addition of the β exponent to the PF metric denominator is a simple yet very effective way to influence the throughput fairness between users because the past throughput of a user can be given more or less weight. Most notably, setting $\beta = 0$ will make the user and resource scheduling oblivious of any past throughputs yielding the behavior of the *max throughput* scheduler introduced in Section 2.2.3.

In our implementation, each TTI n the throughput average $\overline{\gamma}_{k,n}$ is updated using an *exponential smoothing* scheme with smoothing factor $\alpha = 0.001$. The update is based on the previous value $\overline{\gamma}_{k,n-1}$ taking into account if user k was scheduled on PRB f_{idx} (denoted by $x(k, n, f_{\text{idx}}) = 1$) or not (denoted by $x(k, n, f_{\text{idx}}) = 0$):

$$\overline{\gamma}_{k,n} = (1 - \alpha) \cdot \overline{\gamma}_{k,n-1} + \alpha \cdot \left[\sum_{f_{\text{idx}}} x(k, n, f_{\text{idx}}) \cdot r(\text{CQI}(k, f_{\text{idx}}, n)) \right]. \quad (5.2)$$

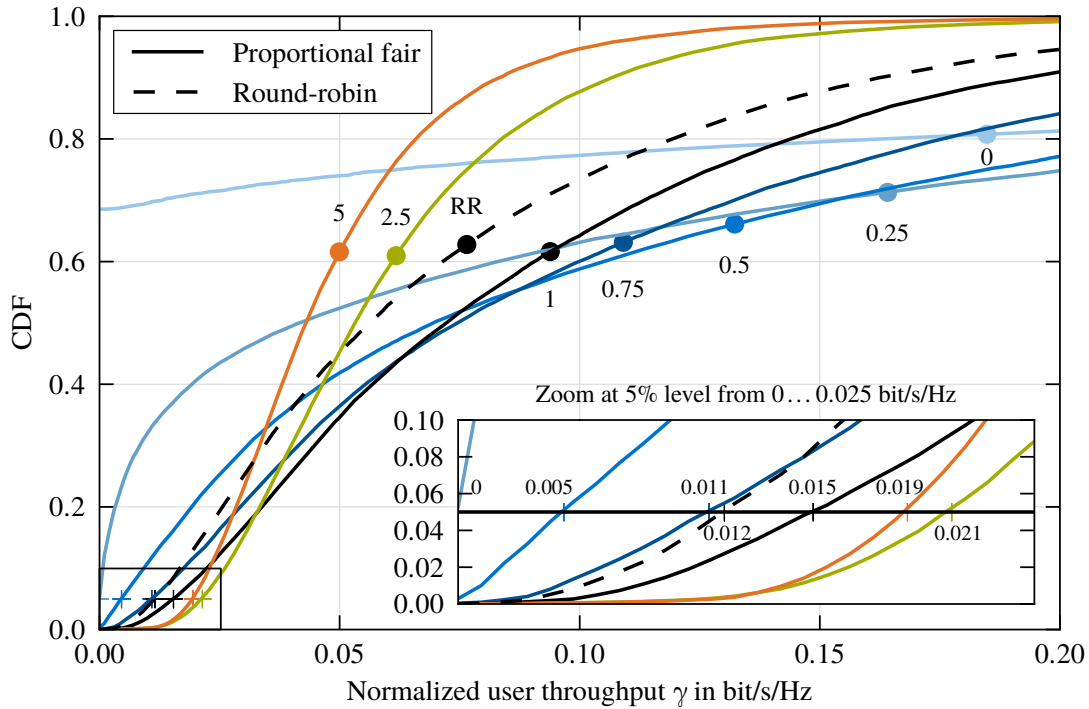
Instead of exponential smoothing, any other scheme could be used to compute past throughput averages. Exponential smoothing, however, is usually chosen due to its simplicity because there is no need to store more than one historic value. The smoothing factor $\alpha = 0.001$ used in IMT_Aphy leads to a rather long memory, which benefits overall system spectral efficiency at the cost of limited short-term fairness.

Fairness Aspects: Cell-Edge User Versus Per-Cell System Spectral Efficiency

The first simulation results presented in Figs. 5.1 and 5.2 examine the influence that the PF history exponent β has on the distribution of throughputs within the cell. All

Table 5.2.: Specific simulation parameters for Figs. 5.1 and 5.2

Parameter	Value
Scenario	urban micro (UMi)
Number of drops	40
Schedulers	round-robin (RR) and proportional fair (PF), PF with exponents $\beta \in \{0, 0.25, 0.5, 0.75, 1, 2.5, 5\}$


 Figure 5.1.: System-wide user throughput distribution achieved by RR (dashed line) and PF scheduling (solid lines) depending on PF history exponents β

simulation parameters listed in Table 5.1 apply to this simulation setup. In addition, the parameters given in Table 5.2 were selected to produce these results.

In the figures, the dashed curve represents the results achieved with the round-robin scheduler, while the solid lines represent the results of the proportional fair scheme with different values for the history exponent β listed in Table 5.2. Figure 5.1 shows the cumulative distribution function (CDF) of the normalized user throughputs γ as defined in (1.2) on page 4. The mean values per setup are indicated with round markers that are labeled with the history exponent β that was used or with “RR” in the case of the round-robin scheduler.

It can be seen that for all curves the throughput distribution is skewed by the higher throughput users because in all cases the mean value is larger than the median value. From the mean per-user normalized throughput we can directly obtain the average per-cell *system spectral efficiency* η , as defined in (1.1), by multiplication with the average

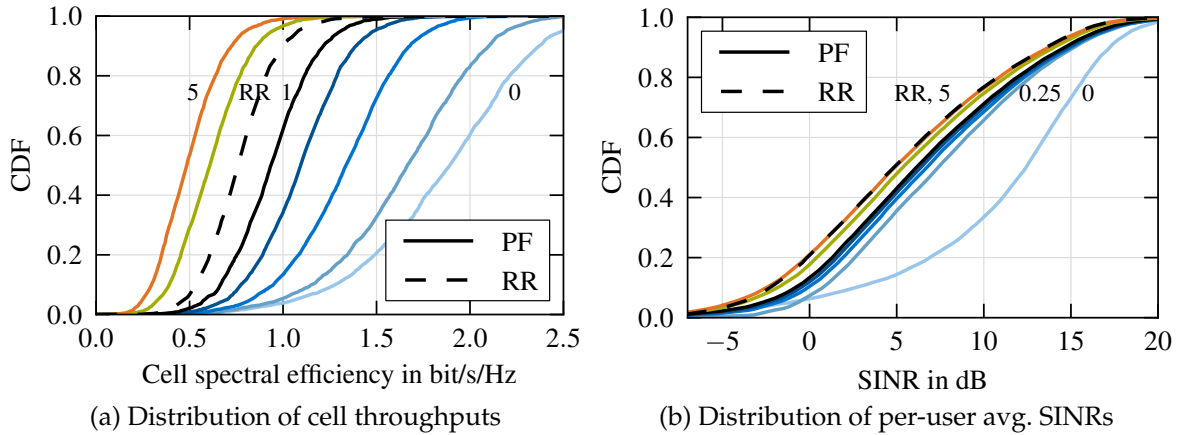


Figure 5.2.: Comparison of cell spectral efficiency and post-scheduling average SINRs per user for RR (dashed lines) and PF (solid lines) schedulers in the UMi scenario; PF schedulers with exponents β between 0 and 5

number of users per cell (10 in this case, see Table 5.1). The other key KPI, that is, the *cell-edge user spectral efficiency* (or *cell-edge user throughput*) is also visible in Fig. 5.1 as the 5% level of the CDF. The inset in the figure shows a zoom around the 5% level and indicates the achieved 5%ile per curve.

For each of the $570 = 10 \cdot 57$ users in a simulation drop, one throughput value is derived at the end of the drop based on the total number of successfully delivered bits. The results from all simulation drops are aggregated into one overall result so that the curves shown in Fig. 5.1 represent the results from $22\,800 = 40 \cdot 570$ individual users. Of course, for each simulation drop new channel realizations are randomly generated as described in Chapter 3, but all simulation jobs of one drop exhibit identical channels making the different runs (e.g., with different schedulers) comparable.

Turning to the results depicted in Fig. 5.1, we can draw two main conclusions. On the one hand, channel-aware scheduling with the basic proportional fair scheduler with $\beta = 1$ outperforms round-robin scheduling over the whole range of user throughputs as the PF CDF lies strictly to the right of the RR CDF. Note that this does not mean that each single user will benefit from PF scheduling, but it does mean that PF scheduling with $\beta = 1$ increases the average spectral efficiency without disadvantaging some users, which is reflected by an also increasing cell-edge user throughput compared to round-robin. On the other hand, we can observe that by tuning the history exponent β , we can achieve a better throughput fairness among users (denoted by a steeper CDF) at the cost of overall spectral efficiency. Alternatively, we can significantly increase the average cell spectral efficiency at the cost of cell-edge users. In the extreme case of maximum throughput scheduling ($\beta = 0$), we can almost triple the spectral efficiency compared to RR, but we leave almost 70% of users with zero throughput.

The vastly different cell spectral efficiencies are also visible in Fig. 5.2a, which depicts the distribution of per-cell spectral efficiencies among the cells in the system where, in this case, $40 \cdot 57$ per-cell results are the basis for the CDF. The figure next to it, Fig. 5.2b,

Table 5.3.: Specific simulation parameters for Fig. 5.3

Parameter	Value
Scenario	urban micro (UMi)
Number of drops	15
Schedulers	round-robin (RR) and proportional fair (PF) with $\beta = 1$
Subband sizes	$N_{\text{PRB}} \in \{50, 25, 16, 12, 10, 8, 7, 6, 5, 4, 3, 2, 1\}$ PRBs per subband

shows the distribution of post-scheduling per-user average SINRs. For each user, one linear SINR average over the whole simulation time of the drop is computed and goes into the distribution that is plotted in the figure. Figure 5.2b helps to understand how the different throughput values per scheduling scheme are achieved. As mentioned above, the channels for a drop are identical between runs with different schedulers and no further Tx or Rx processing is performed that could change the SINRs. The differences between the curves are thus only due to the different probabilities with which good or bad users (in terms of their CQI and thus SINR) are selected for scheduling. The very throughput-fair PF variant with $\beta = 5$ results in SINRs that are only marginally worse than the distribution resulting from round-robin scheduling. In contrast, the maximum throughput scheduler with $\beta = 0$ achieves a significantly better mix of SINRs by serving only the best third of all users as shown in Fig. 5.1.

When we want to evaluate the throughput fairness in the remainder of this thesis, we will refer to the normalized user throughput CDF curve and the cell-edge user throughput (defined as the 5%ile of the normalized user throughput distribution). This way, we will make sure that a cell spectral efficiency increase is not just obtained by an unfair distribution of throughputs between the users in the cell. In the literature, another fairness metric introduced by Jain [JCH84] is often used. Jain's fairness metric maps the user throughput distribution to a dimension-less value in the range from 0 . . . 1. However, we prefer the fairness measures introduced above since they also give an idea of the absolute throughput performance and, in the case of the CDF, give a more complete picture of the throughput distribution.

Frequency-Selective Scheduling Gains

So far, the throughput gains with proportional fair scheduling shown above were achieved with time-domain scheduling only. Though not explicitly stated before, this is a consequence of the wideband CQI feedback, which had been configured according to Table 5.1: without any CQI discrimination between PRBs as well as full-buffer traffic, the PF scheduling metric (5.1), for a given TTI n , always yields the same user $k^*(f_{\text{idx}}, n)$ as the best user, independent of the PRB f_{idx} . Of course, one of the promises connected with channel-adaptive scheduling (see Section 2.2.3) is frequency-selective scheduling that assigns a PRB to a user who is currently expected to have good channel conditions on that PRB. In the following, we will examine the gain from frequency-selective scheduling and discuss the subband-size choice for CQI feedback in LTE.




Figure 5.3.: Scheduling gains in cell spectral efficiency, cell-edge user throughput, and post-scheduling average SINR (right-hand axis) for RR and PF schedulers versus subband size (error bars indicate 95% confidence intervals)

For the evaluation of frequency-selective scheduling gains, we choose simulation settings as summarized in Table 5.3 in addition to those listed in Table 5.1. Now, and in the following, we will always assume a history exponent of $\beta = 1$ when referring to PF scheduling. The main parameter we vary in these simulations is the subband size expressed as the number of PRBs, for which a single CQI value is fed back to the BS. In order to limit the required feedback bandwidth, only a limited number of subbands, whose sizes thus vary with the system bandwidth, can be reported in the LTE standard. For the 10 MHz bandwidth evaluated here, the smallest subband size would be 3 PRBs if the CQI is reported in an aperiodic manner. Conversely, in order to save DL signaling bandwidth, only certain downlink allocation granularities can be signaled. For the sake of this evaluation, we ignore these limitations of the standard and assume subband sizes as indicated in Table 5.3. In the channel-adaptive scheduling process, the per-subband CQI feedback leads to PRB allocations of the same granularity because again the PF metric is identical for all PRBs belonging to the same subband.

In Fig. 5.3 the average normalized user throughputs (solid lines), the cell-edge user throughputs (dashed lines), as well as the average post-scheduling SINRs (dotted lines, with respect to right-hand axis) achieved by PF and RR schedulers are plotted over the indicated range of CQI feedback subband sizes. Here, and in the following, the error bars indicate the two-sided 95% confidence interval. The confidence interval is computed over all simulation drops used to generate the plots. As a general rule, the overall variance and thus the confidence interval shrinks with a greater number of trials (e.g. drops or users) so that we can use the interval size to control the number of simulation drops we generate. We only provide confidence intervals for quantities plotted in linear scale so that, for example, SINRs are not shown with error bars.

The results depicted in Fig. 5.3 indicate that going from wideband CQI feedback (i.e., 50 PRBs per subband) to just a single or very few PRBs per subband, can offer about 40%

Table 5.4.: Specific simulation parameters for Fig. 5.4

Parameter	Value
Scenario	urban micro (UMi)
Number of drops	30 for $N \leq 50$ users per cell, 10 otherwise
Number of MS	on average $N_{\text{MS}} \in \{3, 6, 8, 10, 12, 15, 20, 25, 30, 35, 40, 50, 60, 75\}$ users per cell
CQI subband size	subband (i.e., 3 PRBs per subband) and wideband
Schedulers	round-robin (RR) and proportional fair (PF)

of spectral efficiency gain with PF scheduling. The cell-edge curve shows that there is no trade-off between average and cell-edge user throughput as the cell-edge user throughput increases in a similar way when going to smaller subbands. We also see that the improved throughput is due to an increased average post-scheduling SINR. As expected, the RR scheduler, which already performs worse with wideband CQI feedback, does not benefit from frequency-selective CQI feedback (apart from a small improvement, which can be explained by more accurate link-adaptation).

When comparing the performance of the smallest subband size (1 PRB) to the aperiodic CQI reporting subband size of 3 PRBs per subband, we observe that the performance loss by the larger subband is only marginal compared to the 3-fold savings in required feedback bandwidth. The LTE subband size design choice thus appears to be a very reasonable performance–overhead trade-off.

The frequency-selective scheduling gains presented above are only possible because the BS scheduler can choose between multiple users, which, most likely, have independently fading channels (*multi-user diversity*). With on average 10 users per cell (see Table 5.1), the PF scheduler obviously already has a large-enough multi-user diversity in order to realize sizable gains. How many users are needed on average, and how the gains scale with the number of users per cell, is evaluated next.

Multi-User Diversity Gains

In order to examine the multi-user diversity gains of channel-adaptive scheduling, we conduct a simulation campaign with varying numbers of users per cell. Table 5.4 lists the simulation parameters, and the results obtained with round-robin and proportional fair schedulers are visualized in Fig. 5.4. Since the expected per-user throughput decreases with a growing number of users who have to share the radio resources, we plot the total cell spectral efficiency over the average number of users per cell. As Fig. 5.4 shows, already a small number of users allow the channel-adaptive PF scheduler to obtain a significant advantage over the channel-agnostic round-robin scheduler. While the multi-user gains are smaller and saturate earlier if only wideband CQI feedback is available, the subband CQI feedback (here with the default subband size of 3 PRBs) gives the PF scheduler more freedom to exploit the multi-user diversity leading to a larger gain that only saturates at around 50 users per cell.




Figure 5.4.: Multi-user diversity gains achieved by PF and RR schedulers: cell spectral efficiency versus average number of users per cell in the UMi scenario

Further Aspects of PF Scheduling

After introducing basic properties of channel-adaptive PF scheduling and before evaluating the link adaptation that goes along with scheduling, we will first discuss some further aspects of proportional fair scheduling, including alternative schemes as well as signaling and HARQ retransmission scheduling considerations.

Various variants of PF scheduling have been published in the literature. One popular PF scheme [PPM+07; HT11] separates the scheduling into a so-called time-domain and a frequency-domain step. The idea is to limit the number of users considered for frequency-selective scheduling (step 2) by only admitting a subset of users in the time-domain phase (step 1). Doing so also allows to use separate metrics in the time- and frequency-domain phase. These could be PF metrics for both parts. However, the authors also suggest, for example, to aim at throughput fairness with the time-domain selection, while aiming for a frequency-selective scheduling gain by assigning each user to its preferred parts of the frequency spectrum. An implementation of such a 2-step scheme is available in *IMT_Aphy*. We do not present detailed simulation results here because no significant performance advantage was observed. In addition, with the variation of PF exponents, we already discussed a simple method to trade-off overall system spectral efficiency against cell-edge user throughput. Nevertheless, a 2-step scheduling approach could be interesting for practical deployments because it reduces the computational complexity of the frequency-selective scheduling process. Furthermore, scheduling users less often but with bigger transport block allocations (a consequence of the 2-step approach) could be beneficial for MS power consumption as the receiver could be switched off more often.

In general, the resource scheduling strategy should also be optimized in light of the

carried traffic. For example, web browsing traffic is often bursty so that stretching-out the transmission of one burst (or even of a single IP packet) into many subsequent but small transport block allocations would increase the delay. This can become a problem especially for TCP connections during the *slow-start* phase of the TCP congestion control algorithm [RFC5681]. A higher PHY-layer delay increases the *round trip time* (RTT) of the IP packets, which causes the sender's TCP implementation to perform a slower ramp up of the offered traffic resulting in lower average throughput perceived by the user. In this thesis, we focus on full-buffer scenarios and do not model a TCP/IP stack so that the above-mentioned cross-layer dependencies are not evaluated. Note that IMT_{Aphy} does support system-level simulations with TCP/IP traffic if a suitable traffic generator and TCP/IP stack is configured in the openWNS framework.

Further practical aspects of scheduling in LTE are the limited PDCCH capacity and HARQ retransmissions. Besides UL scheduling grants, the BS also has to indicate the resource allocations (PRBs, MCS, and potentially PMIs and the transmission rank) within the control region resources available for the PDCCH. This can mean that, depending on the DCI formats, the BS might have to restrict the number of users it schedules. It can also dynamically adapt the size of the control region (between 1 and 4 symbols) but doing so changes the rate matching and hence the effective code rates so that some scheduling decisions might have to be updated. In order to reduce complexity, in our simulations we stick to a fixed control region length of 3 symbols and do not check if a given schedule would fit into the available PDCCH resources.

For HARQ retransmissions, a base station even has two scheduling degrees of freedom. First, in the downlink, the BS is free to transmit a HARQ retransmission in any TTI and on arbitrary PRBs. This is different in the uplink, where the retransmission is synchronous, that is, occurs exactly eight TTIs after the preceding failed transmission on the exact same resources. For simplicity reasons, we have also applied the synchronous approach in the downlink for almost all results presented in this thesis even though limited support for asynchronous HARQ operation is available in IMT_{Aphy}. The second degree of freedom for the LTE DL scheduler is the kind of information used for retransmission. Again for simplicity reasons, we only consider Chase combining here and model the retransmission of the initially transmitted bits using the same number of PRBs resulting in the same code rate. The DL scheduler would also be free to perform a CC retransmission with fewer or more PRBs (i.e., a higher or lower effective code rate) or it could transmit additional redundancy bits in an incremental redundancy fashion. For IR, the scheduler can decide which additional bits to transmit (via the redundancy version) and how many bits to transmit. Both aspects are optimization problems in their own right, which we do not consider further here.

5.1.2. Challenges of Outdated CSI

As we have seen above, a channel-adaptive PF scheduler can realize a substantial multi-user diversity gain if quantized CSI in the form of wideband or even subband CQIs is provided by the feedback mechanism. Besides the quantization, another imperfection

Table 5.5.: Specific simulation parameters for Fig. 5.5

Parameter	Value
Scenario	urban micro (UMi)
Number of drops	15
CQI subband size	subband (i.e., 3 PRBs per subband) and wideband
Schedulers	round-robin (RR) and proportional fair (PF)
User speed	$v \in \{0, 1, 2, 3, 4, 5, 6, 8, 10, 12, 15, 17, 20, 25, 30, 35, 45\}$ km/h

in an actual system is always the timeliness of the feedback. In the following, we will thus evaluate, how sensitive the channel-adaptive scheduling gains are to outdated feedback and what kind of link adaptation countermeasures can be taken in order to make the system more robust against channel fluctuations.

According to Table 5.1, we assumed a 6 ms delay between the feedback computation and the time it can be applied by the scheduler. In addition, we modeled an update period of 5 ms so that the feedback is between 6–10 ms old when it is used by the scheduler. While the update period could be reduced to 1 ms in LTE (at the expense of a five-fold signaling overhead), the 6 ms delay is already an optimistic estimate: the MS needs time to generate a report based on measurements, it needs to transmit it, and the BS needs time to decode it and make it available to the scheduler before it can be used. Since the feedback delay values in Table 5.1 are quite representative of an actual system, we will thus vary the temporal fading behavior of the channel by modifying the assumed user movement speeds.

Scheduling with Outdated Channel Feedback

We modify the user speed that is used for the time evolution of the spatial channel model (see Section 3.3.3) and apply values between 0 km/h and 45 km/h, see Table 5.5. Figure 5.5 shows that the PF scheduler can only obtain a substantial multi-user diversity gain at relatively low speeds, whereas the round-robin scheduler only marginally benefits from lower speeds due to a better link adaptation performance.

The coherence time, which characterizes how fast the channel changes making the feedback outdated, is plotted with respect to the right figure axis. According to (2.13) on page 13, it depends on the user's velocity and the system (center) frequency. For the urban micro (UMi) scenario, the center frequency is quite high at 2.5 GHz leading to a short coherence time.

We observe in Fig. 5.5 that the coherence time decreases with higher user velocities and falls below the feedback delays at around 10 km/h. This is also the velocity range at which the gain from adaptive scheduling with subband feedback falls below the performance with wideband feedback. With subband feedback, the greedy PF scheduler tries to schedule its users onto their fading peaks in the frequency-domain. But once the channel changes too fast, the peaks are already gone once the transmission

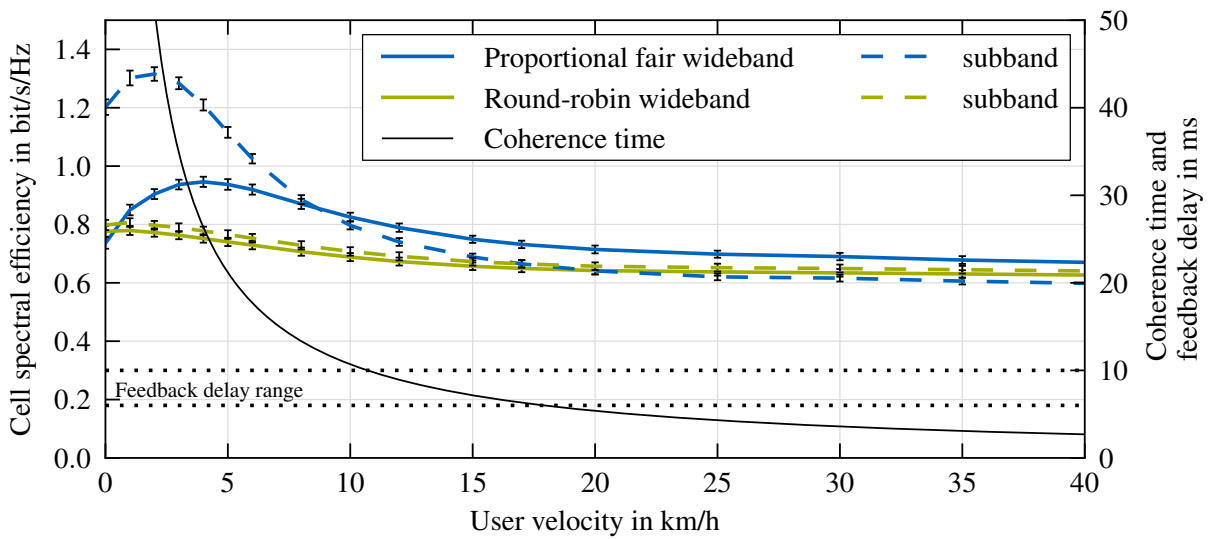


Figure 5.5.: Scheduling gains and coherence time vs. user speed in the UMi scenario

starts. With wideband feedback and wideband allocations (PF in the time-domain), the scheduling is more resilient to the fading.

It is interesting to note that compared to no user velocity (which is unrealistic anyway because even for fixed users, movements of objects in the multipath propagation environment create a slowly fading channel), a slowly changing channel is beneficial because it creates more opportunities for scheduling a user onto a fading peak. The throughput peak in Fig. 5.5 coincides with the default user velocity of 3 km/h in the UMi scenario. In the InH scenario (3 km/h at 3.4 GHz), in the UMA scenario (30 km/h at 2 GHz), and in the RMa scenario (120 km/h at 800 MHz), the default user velocities cause faster fading channels with correspondingly lower expected channel-adaptive scheduling gains, see also the coherence times in Table 3.1 on page 60.

Outdated feedback is a problem that is fundamental to all feedback-based schemes. There have been proposals in the literature to mitigate the performance losses due to feedback aging by predicting future fading behavior [Due07; AVWS12]. While these attempts might be promising in artificial fading channels characterized by fixed Doppler frequencies, real world multi-path propagation environments (e.g., vehicular or pedestrian deployments) often change in unpredictable ways.

Introduction to Outer-Loop Link Adaptation

So far, we examined how the fading channels and outdated feedback pose challenges to the channel-adaptive scheduler, which tries to let its users “ride on the peaks” of their fading channels. With increasing Doppler frequencies and thus less accurate feedback, it becomes harder and harder to hit the peaks. This not only makes the scheduler miss the desired SINR gains, it also complicates the link adaptation as the scheduler tends to overestimate the SINR and thus the possible modulation and coding scheme (MCS) level.

In the previous simulations, we used an unmodified mapping from the reported CQI value to the applied MCS. For this purpose, also LTE systems in the field can use a link-to-system model similar to the one used by our system-level simulator. The reported CQI values per subband correspond to SINR values as shown by the step functions in Fig. 4.5. After converting subband feedback CQI values to per-subband SINR values, and once the BS knows the PRB allocation of a transport block, the BS can compute an expected effective SINR value for the planned transmission by using an SINR compression mechanism (e.g. MIESM) similar to those in the link-to-system model discussed in Chapter 4. This single effective SINR for the scheduled transport block can then be used to lookup a single MCS value for the transport block according to pre-determined MCS-vs-SINR performance curves similar to the BLER curves shown in Fig. 4.4.

One approach to cope with fading channels and potentially outdated feedback is then to subtract a fixed offset Δ from the estimated effective SINR and choose a more robust modulation and coding scheme based on this adjusted SINR. The downside of such a fixed offset is that it is rarely the correct offset. If the offset is too big, the resulting BLER is low but the link adaptation is too conservative and cannot exploit the potential throughput. If the offset is too small, the link adaptation and thus the scheduled rates are too aggressive and the resulting BLER is too high.

In the literature and also based on our experience, a BLER of about 10 % is considered a good trade-off between achieving a high-rate MCS for the initial transmission and risking a retransmission. One very popular link-adaptation strategy called *outer-loop link adaptation* (OLLA) [NAV02; R1-01-0589; PPM+07; WPS+08] thus adapts the estimated SINR (or directly the chosen MCS) by taking the current BLER into account, which is represented by the current ratio of HARQ ACKs and NACKs. Since for each transmitted transport block, that is, up to every 1 ms TTI for a continuously scheduled MS, a HARQ ACK or NACK response is sent, the HARQ feedback can be seen as a second source of frequent (but coarse) CSI feedback besides the CQI.

Even though a HARQ-based outer-loop link adaptation can be considered a standard LTE feature, it is itself not standardized. Thus, we refer to the OLLA mechanism as implemented in `IMTaphy` in the following. In a nutshell, we modify the effective SINR derived from the CQI feedback for a scheduled transport block by adding an offset $\Delta(n, k)$ that is specific to user k and updated each TTI n . For each HARQ ACK feedback, we increment the offset by a small offset Δ_{ACK} (in dB). For each HARQ NACK feedback, we decrement the offset by an offset Δ_{NACK} . We choose the HARQ adjustments Δ_{ACK} and Δ_{NACK} , so that over some long-enough time interval, the expected higher number of small (positive) ACK offsets is compensated by the fewer number of accordingly bigger (negative) NACK offsets. Around the desired operating point, where we aim at a BLER* of 10 %, we expect the following relationship between the numbers N_{ACK} and N_{NACK} of ACKs and NACKs and their respective offsets:

$$0 \stackrel{!}{=} \Delta_{\text{ACK}} \cdot N_{\text{ACK}} + \Delta_{\text{NACK}} \cdot N_{\text{NACK}}, \quad (5.3)$$

where for the target BLER, by the definition of the block error ratio, holds:

$$\text{BLER}^* = \frac{N_{\text{NACK}}}{N_{\text{ACK}} + N_{\text{NACK}}} \quad (5.4)$$

$$\Leftrightarrow N_{\text{ACK}} = \frac{N_{\text{NACK}} - N_{\text{NACK}} \cdot \text{BLER}^*}{\text{BLER}^*}. \quad (5.5)$$

Inserting (5.5) into (5.3) yields the relation (5.6) between the ACK and NACK offsets:

$$\begin{aligned} 0 &= \Delta_{\text{ACK}} \cdot \frac{N_{\text{NACK}} - N_{\text{NACK}} \cdot \text{BLER}^*}{\text{BLER}^*} + \Delta_{\text{NACK}} \cdot N_{\text{NACK}} \\ \Leftrightarrow \Delta_{\text{NACK}} &= -\Delta_{\text{ACK}} \cdot \frac{1 - \text{BLER}^*}{\text{BLER}^*}. \end{aligned} \quad (5.6)$$

In our implementation in `IMTAphy`, we use a target $\text{BLER}^* = 10\%$ and an offset of $\Delta_{\text{NACK}} = 0.03$ dB. The above discussion only considered one spatial layer, but the extension to two or more layers (i.e., two LTE downlink transport blocks) is straightforward. However, independent OLLA processes for single-layer and spatial multiplexing operation could be used to account for possible differences in single- and multi-layer feedback accuracy.

OLLA is not only suitable for adapting the SINR margin under fading scenarios, but it is also a simple method for real-world systems to align different performance models between base and mobile stations. In a way, OLLA also helps the BS to learn how a mobile reports its channel state so that it can compensate a systematic delta between the actual and understood channel state information. Besides different CQI interpretations on the BS and MS side, OLLA also helps to correct the CQI reporting in situations where the MS is blind to certain interference sources. For example, a typical MS estimates the CQI based on the cell-specific reference signals (CRS, see Fig. 2.12). Neighboring cells are usually configured with CRS at different subcarrier positions so that the CRS of the serving cell are interfered by the PDSCH-carrying REs of an interfering cell. If the interfering cell is empty, there will be no interference from the PDSCH of that cell affecting the CRS of the serving cell and hence the MS will report a high CQI value. Some of the data-carrying PDSCH REs of the serving cell, however, will be interfered by the CRS of the interfering cell, which are always broadcast even if no data is transmitted on the PDSCH. This interference is not considered by the CQI reporting, but the OLLA mechanism of the BS will gradually compensate for it.

Performance of Outer-Loop Link Adaptation

In order to evaluate the performance of outer-loop link adaptation, we perform simulations similar to those shown in Fig. 5.5 and vary the user velocity while applying different schedulers with and without OLLA. The parameters are listed in Table 5.6 and the results are depicted in Fig. 5.6. Solid lines denote the performance in the UMi scenario and dashed lines refer to the indoor hotspot (InH) scenario. The UMi results in Fig. 5.6a for the RR and PF schedulers without OLLA are comparable to the results plotted in Fig. 5.5 with the difference that, due to 25 instead of 10 users per cell, slightly higher multi-user diversity gains are achieved here. In the UMi scenario, the effect of OLLA is visible for speeds higher than 5 km/h, where OLLA prevents that the performance of the PF scheduler is degraded by an increasingly overestimated SINR.

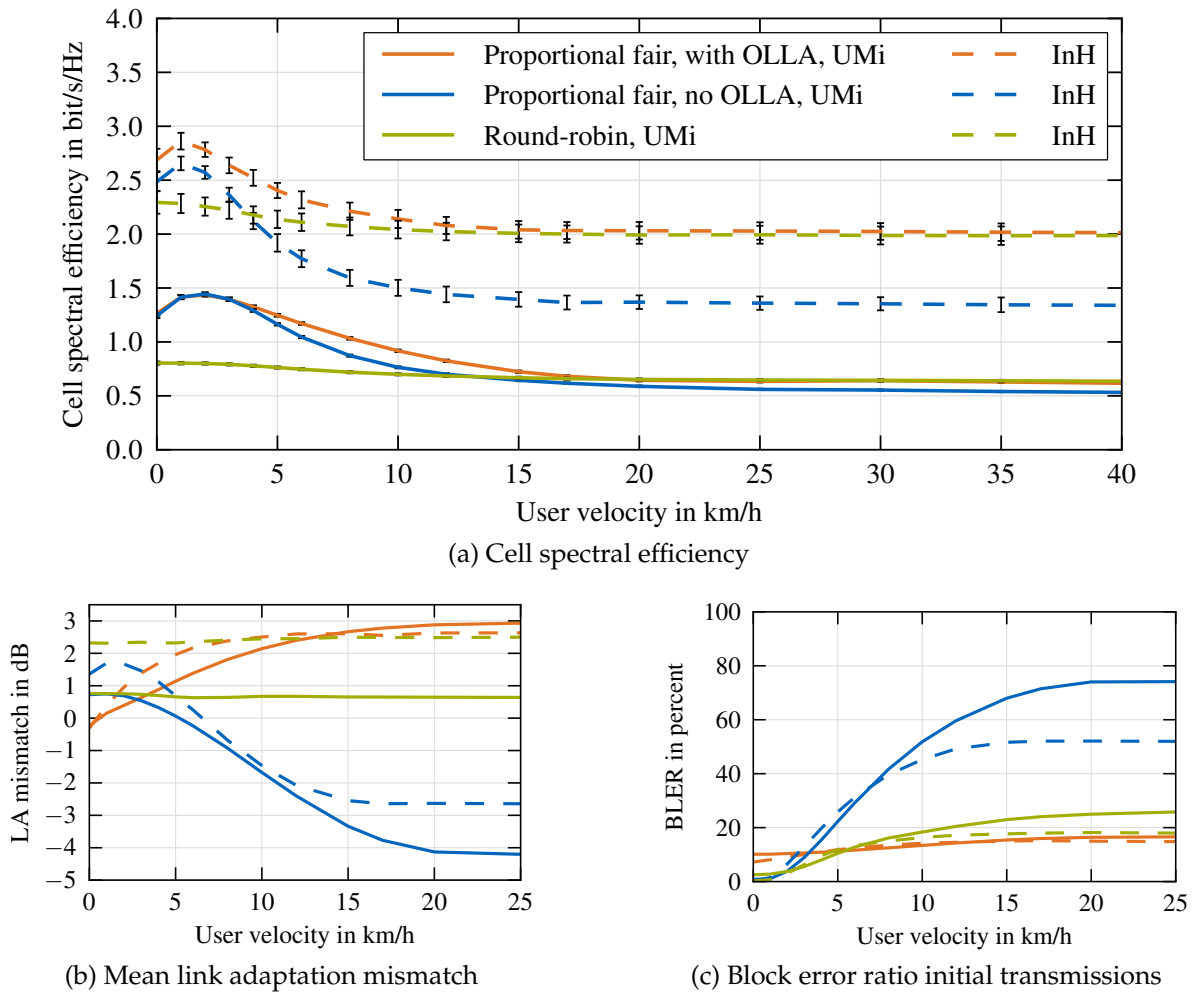


Figure 5.6.: Link adaptation mismatches (actual – estimated SINR), resulting BLERs, and spectral efficiencies for RR and PF schedulers with and without OLLA in the UMi and InH scenarios versus user velocities (legends for Figs. (b) and (c) see Fig. (a))

Table 5.6.: Specific simulation parameters for Figs. 5.6 and 5.7

Parameter	Value
Scenario	urban micro (UMi) and indoor hotspot (InH)
Number of drops	15
Number of MS	on average $N_{MS} = 25$ users per cell
User speed	$v \in \{0, 1, 2, 3, 4, 5, 6, 8, 10, 12, 15, 17, 20, 25, 30, 35, 45\}$ km/h for Fig. 5.6 and $v \in \{10, 30\}$ km/h for Fig. 5.7
Link adaptation	Fixed offset of $\Delta(n, k) = 0$ dB and outer-loop link adaptation (OLLA) with target $BLER^* = 10\%$ and $\Delta_{NACK} = 0.03$ dB
CQI subband size	subband (i.e., 3 PRBs per subband)
Schedulers	round-robin (RR) and proportional fair (PF)

Without OLLA, the PF scheduler even performs worse than RR scheduling for higher user velocities. The same effects are visible in a more drastic way in the InH scenario (dashed lines): without an OLLA correction, the greedy frequency-selective PF scheduling leads to an overestimation, which results in PF throughputs much worse than the RR performance, as soon as the velocities increase to more than just a few km/h. The reason why the indoor hotspot scenario creates a more challenging environment is that, with just two BSs in the scenario, a fading dip of the interfering cell can create a too short-lived SINR peak for the greedy PF scheduler to successfully exploit. Also, bursty traffic in interfering cells can create SINR fluctuations that are unpredictable because the scheduling can change every 1 ms TTI. Since we only consider full-buffer traffic in this thesis, we do not evaluate such situations.

The overestimation of the expected SINR in the setups discussed above is shown in Fig. 5.6b. There, we plot the average difference between the effective SINR with which a transport block (spanning multiple PRBs) is actually received and the effective SINR that had been estimated for that transport block based on the previously reported CQIs. This so-called *link adaptation mismatch* is quite fading-independent for the RR scheduler because the 10 MHz-wide (50 PRBs) allocations average out the frequency-selective fading. Note that there is a systematic positive mismatch which can be explained by the CQI granularity. The CQIs (see Fig. 4.5a on page 102) cover a dynamic range of about 30 dB in roughly 2 dB steps so that the CQI feedback implicitly has a certain safety margin. In the InH scenario, the RR scheduler exhibits an even larger positive LA mismatch because in the InH scenario the SINRs are often much larger (see Fig. 4.6a on page 106) than necessary for the highest CQI and MCS.

What is more interesting for our discussion of OLLA performance is the increasingly negative LA mismatch experienced by the PF scheduler without OLLA for higher speeds. As shown in Fig. 5.6c, the increasingly negative LA mismatches resulting from an overestimation of SINRs lead to high block error ratios above 50 % for initial transmissions once the speed increases to more than 10 km/h. Due to the overall lower SINR range in the UMi scenario (see Fig. 4.6a), the BLERs are even higher than in the InH scenario, which is also reflected by a higher relative throughput loss between the highest and lowest points in the throughput curve in Fig. 5.6a. With OLLA, in contrast, the BLER stays relatively constant as the user velocity is varied in Fig. 5.6c. Note that for low speeds, the OLLA mechanism tunes the link adaptation to be more aggressive, (i.e., $\Delta(n, k) > 0$ offsetting the implicit CQI safety margin) leading to a BLER around the 10 % target. For higher speeds, the BLER rises slightly above the 10 % target in our simulations, which indicates that, for the small OLLA offset $\Delta_{\text{NACK}} = 0.03$ dB, the 250 ms settling time is a bit too short.

Based on the same simulations, in Fig. 5.7 we take a closer look at the distribution of link adaptation mismatches with different schedulers with and without OLLA for two exemplary speeds at 10 km/h (denoted by solid lines) and 30 km/h (denoted by dashed lines). In Fig. 5.7 we plot the distribution of the individual mismatches from each and every transport block transmitted during the whole simulation considering all users. The higher velocity leads to a wider variance of mismatches. For the RR scheduler, the mismatches are symmetrically distributed around the small positive




Figure 5.7.: Comparison of link adaptation mismatch between RR and PF scheduling with and without OLLA for 10 km/h and 30 km/h

average mismatch already shown in Fig. 5.6b. The PF scheduler without OLLA exhibits a significant tail towards negative mismatches resulting from overestimation, which occurs for more than 60–80% of all transmissions. The activation of OLLA shifts the mismatch CDFs as it adapts the SINR offset $\Delta(n, k)$ according to the current BLER. The advantage of the presented OLLA scheme is that the offsets are automatically adapted without the need for the BS to know the cause (e.g., the user’s velocity) of a user-specific elevated BLER. We will always use the OLLA mechanism in the subsequent simulation results. Once we introduce spatial processing at the BS or MS, we will observe further sources of link adaptation mismatches besides user velocities.

5.2. Spatial Processing at the Receiver (SIMO)

In this section, we now extend the scope of RRM schemes by adding spatial processing at the receiver to the picture. Since we focus on the downlink, the MS is the receiver. The LTE standard [36.101] requires all LTE mobiles to be equipped with at least two antennas because substantial gains can be obtained with more than one antenna as we will present in this section. There is no theoretical upper limit to the number of Rx antennas that can be employed. The limitation, however, comes from cost and space constraints because additional Rx antennas need to be properly separated in space (or feature different polarization or directional antenna patterns) in order to provide a gain with spatial processing.

In the first part of this section, we will thus evaluate how the gains scale with the number of available Rx antennas. Of course, the performance also depends on how the signals from the multiple antennas are combined. Here, we will evaluate the MRC and IRC linear receive filters introduced in Sections 2.1.2, 2.2.2 and 4.2.2. In the second part of this section, we will evaluate how channel and IPN covariance matrix estimation errors negatively affect the performance of the Rx spatial processing schemes. Still, all

Table 5.7.: Specific simulation parameters for Fig. 5.8 and Fig. A.1

Parameter	Value
Scenario	urban micro (UMi) for Fig. 5.8 and UMa for Fig. A.1
Number of drops	5
Number of MS	on average $N_{MS} = 10$ users per cell
BS antennas	1 vertically polarized antenna
MS antennas	array of $U \in \{1, 2, 3, 4, 5, 6, 7, 8\}$ vertically polarized antennas with $\lambda/2$ spacing
Link adaptation	with and without outer-loop link adaptation (OLLA)
CQI subband size	subband (i.e., 3 PRBs per subband)
Scheduler	proportional fair (PF)
Channel estimation	perfect estimation of the serving channel and of the IPN covariance matrix

simulations in this section assume a single Tx antenna at the BS (i.e., TM1) so that only the transmission of a single spatial layer is possible.

5.2.1. Performance Gains from Multiple Rx Antennas

As of 2014, there is hardly any mobile phone or tablet on the market with more than two receive antennas for cellular communications, and also MIMO wireless local area network (WLAN) routers usually only feature up to three antennas. However, more than two Rx antennas should be feasible if cross-polarized antennas are used, if higher frequency bands are used (e.g. around 5 GHz) and, of course, for larger tablets or notebooks. Independent of the current situation—and as a motivation for future improvements—we examine the performance gains when increasing the number of Rx antennas from one to eight. These and the remaining simulation parameters are summarized in Table 5.7, which again extends Table 5.1.

Figure 5.8a shows the per-cell spectral efficiency (top) as well as the cell-edge user throughput (bottom) in the UMi scenario achieved with the MRC and IRC receive filters depending on the number of MS antennas. The performance with one antenna is identical to the previously discussed results (see, e.g., Fig. 5.5 for the UMi default speed of 3 km/h). Additional antennas can significantly increase the throughput, especially if the IRC filter is used. This is true for the average system spectral efficiency, which roughly doubles when going from one Rx antenna to four Rx antennas with IRC filtering. The gain is even more pronounced for the so-called cell-edge users since these usually experience high interference levels where the IRC filter can be most helpful. The IRC cell-edge user throughput almost scales linearly with the number of Rx antennas and more than doubles when going from one to two antennas. The MRC receiver also allows significant, but comparably lower, gains in terms of cell-edge user throughput. For MRC and up to 2 Rx antennas with the IRC filter, no outer-loop link adaptation is necessary at the low default speeds in the UMi scenario. However, the

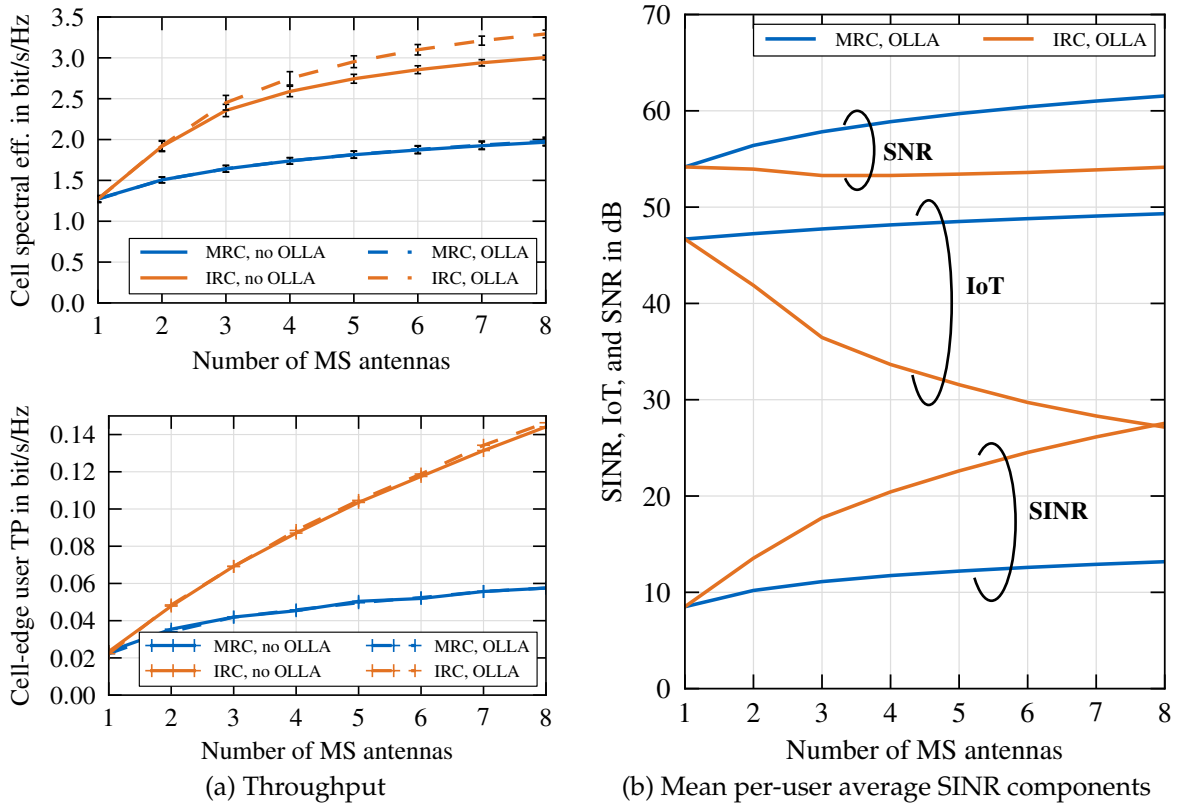


Figure 5.8.: Throughput (TP) gains in terms of system spectral efficiency (top left) and cell-edge user throughput (bottom left) depending on the number of Rx antennas, receive filter, and link adaptation scheme in the UMi scenario; post-receiver SINRs separated into mean SNR and mean IoT (on the right)

spectral efficiency with the IRC receiver with more than 2 Rx antennas further benefits from the use of OLLA, which adds another OLLA use case in addition to higher user velocities. In the Appendix on page 213, we show the equivalent figure derived in the UMa scenario. There, OLLA is also beneficial for the MRC receiver.

Turning to the other part of the figure, Fig. 5.8b allows to understand how the MRC and IRC filters obtain the presented throughput gains. The post-combining SINR produced by the MRC and IRC filter grows with the number of Rx antennas but shows a much higher gain for the IRC receiver that directly translates into the throughput advantage of the IRC filter as visible in Fig. 5.8a. The other two curves, the *signal to noise ratio* (SNR) and *interference over thermal* (IoT), allow to separately analyze the nominator (i.e., the signal power) and the denominator (i.e., the IPN power) of the plotted SINR. Since we consider post-processing values, it does not make sense to evaluate absolute desired signal or interference powers because these could be arbitrarily scaled by the Rx filter. By relating the filtered desired signal power and the filtered interference to the also filtered thermal noise power, we obtain scaling-invariant metrics that allow us to judge the net effect of the filter on the signal and interference powers.




Figure 5.9.: Normalized polar plots in (linear scale) of 4 Rx antennas rank-1 receive filters for different IPN covariance matrix estimation schemes

With respect to the MRC filter, we observe in Fig. 5.8b that by combining the signals received on the up to eight Rx antennas, the MS can achieve a receive beamforming SNR gain (see equations (2.22) and (2.23) on page 17) of about 7.5 dB compared to a single Rx antenna. At the same time, the MRC filter collects only about 2.5 dB of additional interference power (i.e., IoT) while beamforming towards its desired transmitter, yielding a net SINR gain of about 5 dB in total. Note that the increase in IoT might seem surprising at first glance because the MRC-based Rx beamforming leads to a kind of spatial filtering. The resulting beam pattern mostly collects interfering signals with a spatial signature that is similar to that of the desired signal and suppresses other interference sources as illustrated in Fig. 5.9 (see the IRC and MRC curves; the others will be discussed later). In this case, there are two main reasons why we observe an increase in IoT. First, the dense UMi scenario with wraparound modeling surrounds the users with interferers so that an Rx beamformer tailored for the desired transmitter still captures and amplifies a number of interferers (among them the two co-located cell sectors). Second, the desired signal in a multipath propagation environment might arrive with a high angular spread so that the Rx beam needs to cover a larger angle collecting more interference.

The IRC filter, as the name *interference rejection combining* suggests, achieves its significant SINR gain by rejecting a big part of the interference as visible in the IRC IoT curve. The more Rx antennas the receiver has available, the higher the number of interferers (or significant interference multipaths) the IRC filter can suppress. In the UMi scenario shown in Fig. 5.8b, the IoT decreases by almost 20 dB when using eight antennas instead of a single one. At the same time, the purpose of the IRC filter to suppress the interference conflicts with the goal to achieve a combining gain with respect to the desired transmitter (see also the illustration in Fig. 5.9 where the IRC beams do not directly point toward the direction of the desired signal). As we see in Fig. 5.8b, the IRC SNR stays roughly constant when varying the Rx antennas or even decreases slightly. Note that for more than 4 Rx antennas, the increase of the IRC SINR towards 30 dB pro-

duces diminishing throughput gains because many links are limited by the maximum efficiency of the highest 64-QAM MCS in LTE. More Rx antennas should thus either be used for spatial multiplexing, or higher modulations should be added if a high number of Rx antennas becomes feasible, for example, due to higher frequencies.

To summarize the observations of our SIMO simulations, we have seen that in terms of throughput, every additional antenna (that can be fit onto the device with sufficient spacing from the other antennas) contributes a significant gain, especially at the cell edge. To fully utilize the Rx antennas in a multi-cell scenario, the MS has to perform interference rejection combining, which means it has to estimate the IPN covariance matrix. The relative performance gain is largest when going from one to two antennas so that the LTE requirement of at least 2 Rx antennas seems reasonable. Note, however, that operating with a single receive antenna is—in principal—possible for any transmission scheme that just transmits one spatial layer. Just using a single Rx antenna can even be beneficial to the user if the MS can switch off the second receiver chain and use less complex receiver processing in order to save energy.

5.2.2. Serving Channel and IPN Estimation Impairments

In the evaluation of spatial processing gains at the receiver, we assumed perfect channel and IPN covariance matrix estimation. As already discussed in Section 4.3, both the serving channel matrix \hat{H} (needed for MRC and IRC filtering) and the interference plus noise covariance matrix \hat{R}_{IPN} (needed for IRC only) have to be estimated by the MS. In the following, we will evaluate the performance impacts of estimation errors that are modeled as described in Sections 4.3.1 and 4.3.2.

Channel Matrix Estimation

According to our channel matrix estimation error model defined in equations (4.19) and (4.20), the additive error term is a random matrix, whose covariance matrix is the scaled IPN covariance matrix. The processing gain $1/G_{\text{est}}$, by which the error term is scaled down compared to the IPN covariance, depends on the number of reference signals available for channel estimation. In the simulations presented here, we vary the processing gain G_{est} from 0 dB (corresponding to just one CRS per Tx antenna port) to 12 dB, which corresponds to 16 CRS available for channel estimation per Tx antenna port (the latter implying that the channel coherence allows using also CRS from adjacent PRBs and/or TTIs for channel matrix estimation purposes). Table 5.8 lists all simulation parameters of interest for the simulations presented in the following. Figure 5.10 provides results for two and four receive antennas in Fig. 5.10a and Fig. 5.10b, respectively. We plot the CDF of the instantaneous SINR, which means that the SINRs of each PRB belonging to each and every transport block transmitted during the whole simulation time are considered in the distribution without any per-user averaging. The solid curves represent the SINRs obtained from MRC filtering whereas

Table 5.8.: Specific simulation parameters for Fig. 5.10

Parameter	Value
Scenario	urban micro (UMi)
Number of drops	10
BS antennas	1 purely vertically polarized antenna
MS antennas	2 and 4 purely vertically polarized antennas with $\lambda/2$ spacing
Receive filter	MRC and IRC with perfect IPN covariance estimation
Scheduler	proportional fair (PF)
Channel estimation	perfect and IPN-covariance-colored (see (4.20) in Section 4.3) with processing gains $G_{\text{est}} \in \{0, 1, 3, 6, 12\}$

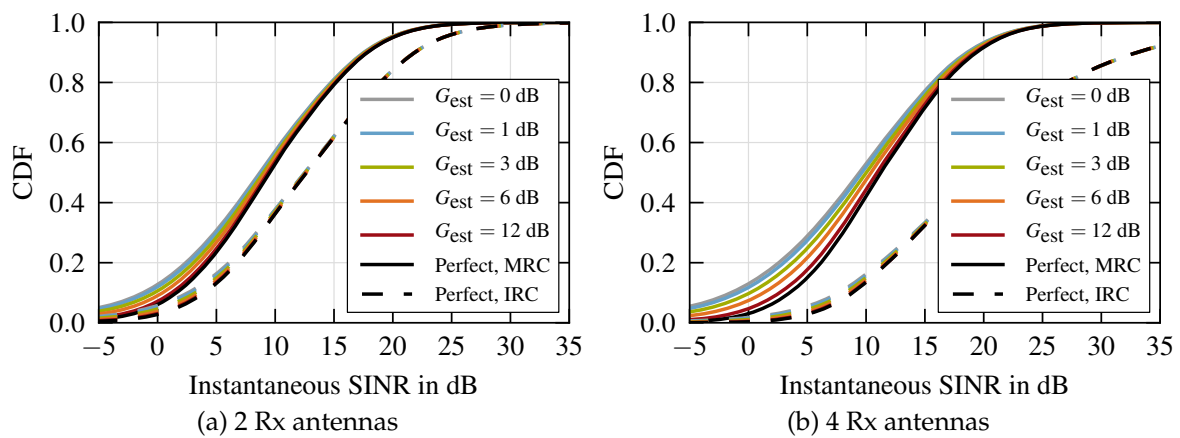


Figure 5.10.: Instantaneous SINR distribution for different serving channel estimation processing gains G_{est} compared to perfect serving channel estimation for 2 and 4 Rx antennas with MRC and IRC receivers in the UMi scenario

the dashed curves represent the SINRs obtained from IRC filtering with perfect IPN covariance matrix estimation.

For each receiver type, the best results are, obviously, achieved by perfect serving channel matrix estimation (plotted in black). As expected, the deviation from the result with perfect estimation grows with a smaller processing gain. We observe that the IRC receiver is not much affected by estimation errors regarding the serving channel because, as seen in Fig. 5.8b, it primarily adapts to the interference, and not the serving channel. This is in contrast to the MRC receiver, which only adapts to the serving channel. Accordingly, serving channel matrix estimation errors have a larger impact on the instantaneous SINRs with the MRC receiver, especially for lower SINRs. The losses are bigger for 4 Rx antennas compared to 2 Rx antennas. With 4 Rx antennas, they can reach more than 5 dB when comparing the performance of perfect serving channel estimation to the error model with a processing gain of only 0 dB.

Despite possible SINR degradations with imperfect serving channel matrix estimation, we will always assume perfect estimation in the remainder of the thesis (as we did

Table 5.9.: Specific simulation parameters for Figs. 5.11, 5.12, A.2 and A.3

Parameter	Value
Scenario	urban micro (UMi) in Figs. A.2 and A.3 and urban macro (UMa) in Figs. 5.11 and 5.12
Number of drops	10
BS antennas	1 vertically polarized antenna
MS antennas	2 and 4 vertically polarized antennas with $\lambda/2$ spacing
Schedulers	proportional fair (PF)
Receiver filter	MMSE filter including IPN covariance matrix estimate
IPN covariance estimation	perfect, diagonal matrix, scaled identity, and Wishart model with $N_{\text{Sample}} \in \{4, 6, 12, 16, 36, 132\}$

previously). Since we can assume that a large-enough number of CRS is usually available for estimating the channel matrix, a high processing gain can be achieved that only leads to small deviations from the result with perfectly-known channel matrices. Besides these SINR losses from filter weight mismatches, there is also a *demodulation loss* in the inner receiver because the phase and amplitude are not perfectly estimated. This additional loss also depends on the processing gain and can dominate the combining loss discussed above [MF11]. Note that we also do not further consider such a demodulation loss in the following.

IPN Covariance Matrix Estimation

In addition to the channel matrix estimation error model, we also introduced an error model for the IPN covariance matrix estimation in Section 4.3.2. Here, we will discuss performance evaluation simulations conducted with simulation parameters as listed in Table 5.9. We compare the effect IRC filtering with different IPN covariance matrix accuracy levels has on the post-receiver SINR and the cell-edge user and system spectral efficiency. Four different IPN covariance matrix estimation types are considered:

1. a *perfect* IPN covariance matrix $\hat{\mathbf{R}}_{\text{IPN}} = \mathbf{R}_{\text{IPN}}$,
2. a *diagonal* matrix with $\hat{\mathbf{R}}_{\text{IPN}} = \text{diag}(\text{diag}(\mathbf{R}_{\text{IPN}}))$,
3. a *scaled identity* matrix with, $\hat{\mathbf{R}}_{\text{IPN}} = \text{tr}(\mathbf{R}_{\text{IPN}}) / N_{\text{Rx}} \cdot \mathbf{I}_{N_{\text{Rx}}}$
4. the *Wishart*-matrix-based model introduced in Section 4.3.2 where the estimation accuracy depends on the number N_{Sample} of samples, see (4.26) on page 98.

As outlined in Section 4.3.2, the number of useful samples in an LTE system depends on the number of CRS REs if the IPN covariance is to be estimated on those REs for which the contribution from the serving BS is known and can thus be canceled. Since the interference situation can be different on adjacent PRBs and TTIs, only one single PRB can be used for the estimation. The minimum number of CRS outside the control region and thus the minimum number of samples is $N_{\text{Sample}} = 6$ in the case of one Tx antenna, see Fig. 2.12. In addition, we choose other reasonable values: $N_{\text{Sample}} = 12$




Figure 5.11.: Instantaneous SINR achieved by the MMSE(-IRC) receiver using different IPN covariance estimation methods in the UMa scenario

(2 Tx ports), $N_{\text{Sample}} = 16$ (4 Tx ports), $N_{\text{Sample}} = 36$ (2 CRS ports and DM-RS for more than two layers), and $N_{\text{Sample}} = 132$ (if the IPN covariance is estimated after detecting and canceling the PDSCH REs outside the control region). Together with $N_{\text{Sample}} = 4$, which is the minimum number of samples necessary to obtain an invertible 4×4 matrix, these are the parameter choices used in our evaluation of the Wishart-based IPN covariance error model, see Table 5.9.

First, we discuss the IPN covariance estimation error impact on the instantaneous SINR after IRC Rx filtering as shown in Fig. 5.11. As expected, the scaled identity (no interference rejection at all) and the perfect IPN covariance estimation are the two extremes. The SINR gains achieved with IRC combining based on imperfect IPN covariance matrices grows with the estimation accuracy as reflected by the assumed number of samples in the Wishart model. In the two Rx antenna case, already four samples allow a consistent SINR gain that hardly differs from the perfect case as soon as 16 or more samples are available. With four Rx antennas, four samples are not always enough to benefit from IRC combining. For example, for the low SINR regions (i.e., for cell-edge users, see also Fig. 5.12) the estimation error with only four samples is so large that a scaled identity IPN covariance matrix estimate \hat{R}_{IPN} , and thus no interference suppression, offers better performance. But with more samples, the performance with the Wishart model quickly converges to the perfect IPN estimation performance. We can conclude that the impressive SIMO gains presented in the previous section, which assumed IRC filtering based on perfect IPN covariance estimation, appear to be achievable with the limited amount of samples that are available in one LTE PRB / TTI for IPN estimation purposes.

The scatter plot in Fig. 5.12 illustrates the same situation in terms of spectral efficiency. It demonstrates that every additional level of IPN covariance matrix estimation accuracy improves both the cell-edge user throughput (y-axis) as well as the system spectral efficiency (x-axis). Since the potential for interference reduction is higher at the cell edge, the relative gains with respect to cell-edge user throughput are higher




Figure 5.12.: Cell spectral efficiency and normalized cell-edge user throughput achieved by MMSE(-IRC) linear receivers with different IPN covariance matrix estimation models in the UMa scenario

than those for the system spectral efficiency. In Figs. 5.11 and 5.12 we considered the performance in the UMa scenario but the situation is very similar in the UMi scenario, for which corresponding figures can be found in the Appendix (Figs. A.2 and A.3).

Overall, the evaluation of channel estimation and IPN covariance matrix estimation possibilities presented in this section suggest that (at least in the SIMO case discussed here) only small system-level performance losses due to inaccurate channel or IPN estimation are to be expected. The small difference between the Wishart-based error model and the perfect IPN estimation is also visualized in Fig. 5.9 where an exemplary Rx beam pattern resulting from perfect IPN estimation and the Wishart-based model with $N_{\text{Sample}} = 4$ is plotted. As a reference, the patterns from the interference covariance-agnostic diagonal and scaled identity models are also shown, which are also similar to each other. As a consequence of the small differences between the Wishart-based error model and the perfect IPN covariance estimation, we will mostly assume perfect IPN estimation in addition to perfect serving channel estimation in the remainder of the thesis in order to simplify the simulation assumptions.

5.3. SFBC Transmit Diversity in LTE

The first multi-transmit-antenna scheme we want to consider is transmit diversity with SFBC as introduced in Section 4.2.3. It is available in LTE as a dedicated transmission mode (TM2), as a fallback scheme in each other transmission mode except TM1, and it is used for PDCCH transmissions if more than one Tx antenna is available at the BS. We will first demonstrate the benefits of SFBC in a context that is representative for

Table 5.10.: Specific simulation parameters for Fig. 5.13 and Table 5.11

Parameter	Value
Scenario	urban macro (UMa)
Number of drops	5
Number of MS	on average $N_{\text{MS}} = 25$ users per cell
BS antennas	2 vertically polarized antennas
MS antennas	array of $U \in \{1, 2, 4\}$ vertically polarized antennas with $\lambda/2$ spacing
Receiver	MRC filter based on perfect channel estimation
Transmission mode	TM1 and TM2
Link adaptation	without outer-loop link adaptation (OLLA)
Schedulers	frequency-domain round-robin, allocating only 1 PRB to each user
Bandwidth	25 PRBs, i.e., 5 MHz




Figure 5.13.: Comparison of instantaneous SINR and per-user normalized throughput between single BS Tx antenna (TM1) and 2 Tx antenna SFBC transmit diversity operation (TM2) with special 1 PRB per user round-robin scheduling for 1, 2, and 4 Rx antennas at the MS

low-rate-PDCCH and non-adaptively-scheduled RRC control signaling transmissions. After that, by showing the results of a small system-level simulation campaign, we will motivate why transmit diversity is not a good choice for high system spectral efficiency if channel-adaptive scheduling is an alternative.

5.3.1. Higher Robustness for Unscheduled Low-Rate Transmissions

We demonstrate the key advantage of SFBC, namely increased robustness against fading on one of the antenna pair propagation paths, by considering a corner case. We use a variant of the round-robin scheduler that performs the round-robin scheduling in the frequency-domain (similar to the RR variant discussed for uplink calibration purposes in Section 4.5.2) by allocating a single PRB to each user in each TTI. That way,

Table 5.11.: Initial transmission BLER performance of TM1 and TM2 with special 1 PRB per user round-robin scheduling

TM	1 Rx	2 Rx	4 Rx
1	37.6 %	26.3 %	19.8 %
2	32.2 %	21.5 %	15.1 %

we allocate the smallest possible physical transmission resource for the PDSCH in a channel-agnostic way. This can be representative for small downlink RRC packets that are sent infrequently when the MS is not otherwise scheduled and thus only rarely provides CSI feedback. It is also the closest equivalent to a PDCCH transmission we can create with our system-level simulator.

Table 5.10 lists the simulation parameters specific to the used setup. In Fig. 5.13 we show two key metrics from the simulation results where we compare single antenna transmissions with TM1 (dashed lines) against TM2 (solid lines) for different numbers of Rx antennas that are available to the MRC filter. The instantaneous SINR CDF in Fig. 5.13a shows an SINR gain of up to 2 dB for low SINRs if only a single Rx antenna is used. This gain can be explained by the SFBC diversity combining that avoids deep fades on one antenna pair by always yielding the average gain of all Tx–Rx antenna pairs as shown in Section 4.2.3. Accordingly, TM2 has slightly smaller SINRs in the high SINR region because also fading gains on one Tx–Rx pair are offset by averaging with the other antenna pairs. The gain vanishes as more Rx antennas are used because MRC combining provides a similar diversity effect in addition to the combining gain. Overall, we see in Fig. 5.13b that, with the considered special channel-agnostic frequency-domain RR scheduling, a small throughput gain is realized with the SFBC scheme. The diversity gains are also visible in Table 5.11, which provides the average BLERs for the considered setups. Note that overall the BLERs are high because at the default speed of 30 km/h, the UMa channel is fading fast and no OLLA scheme is used. In addition, with just a single PRB per transport block, there is not much inherent frequency-diversity that the Turbo decoder could otherwise exploit to successfully recover the transport block if just a few PRBs experience significant fading drops.

5.3.2. Disadvantages of SFBC with Channel-Adaptive Scheduling

The use case above, for which we could show small gains of SFBC transmission, is not representative for high throughput transmission setups where usually many more PRBs per transport block are scheduled in a channel-adaptive way. In fact, the averaging effect of transmit diversity schemes leading to less SINR variability is detrimental to the multi-user diversity gains achieved by channel-adaptive scheduling, which in contrast requires a high degree of variability in the channel.

We conduct a small simulation campaign to evaluate the performance of TM2 in comparison to TM1 in an urban micro scenario with channel-adaptive PF scheduling using

Table 5.12.: Specific simulation parameters for Fig. 5.14

Parameter	Value
Scenario	urban micro (UMi)
Number of drops	5
BS antennas	2 purely vertically polarized antennas with $\lambda/2$ spacing
MS antennas	2 purely vertically polarized antennas with $\lambda/2$ spacing
Transmission mode	TM1 and TM2
Receivers	MRC and IRC filters with perfect channel and IPN covariance matrix estimation
Link adaptation	fixed link adaptation (no OLLA)
Schedulers	proportional fair (PF)

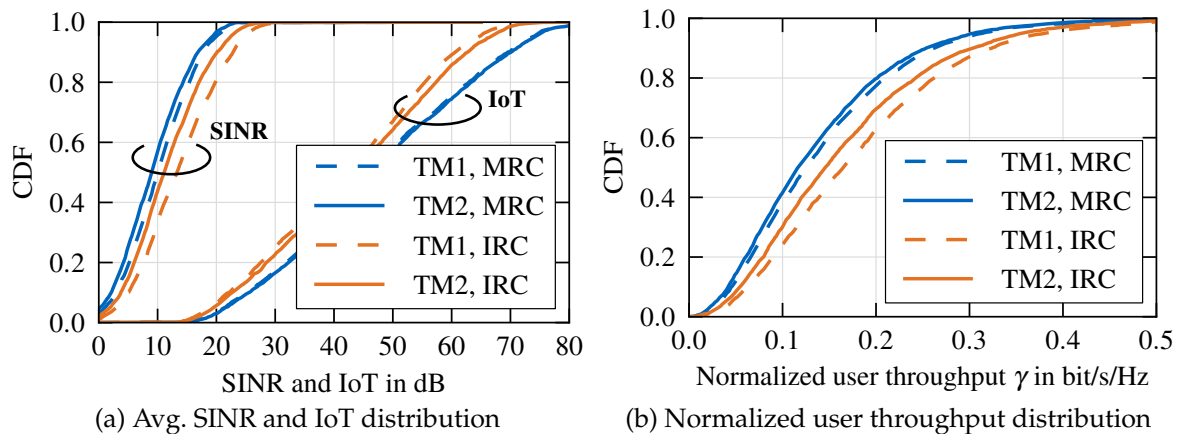


Figure 5.14.: System-level performance of SFBC with channel-adaptive scheduling

a 2×2 antenna configuration. Note that while TM1 only uses a single Tx antenna, we still model a 2 Tx antenna CRS overhead for a fair comparison. Table 5.12 lists the remaining simulation parameters. The performance results in terms of long-term average SINR and IoT as well as the resulting normalized user throughput distributions are shown in Fig. 5.14. Obviously, in this setup the performance of SFBC with TM2 is inferior to a simple single-antenna transmission using TM1 as visible in Fig. 5.14b. Besides expected lower multi-user diversity gains, Fig. 5.14a reveals another factor limiting the performance of SFBC. We observe that a part of the lower average SINRs with TM2 can be explained by a higher interference over thermal ratio, especially for the IRC filter. The reason for this is that an SFBC transmission in a neighboring cell appears as a source of two-layer interference to the receiver (see Section 4.2.3). With effectively twice the amount of interferers (although at identical total power), the IRC filter faces more challenges to suppress TM2 inter-cell interference resulting in higher IoTs, lower SINRs, and thus lower throughputs.

The disadvantages of transmit diversity are representative of a general trend: The use of CDMA-based spreading in UMTS, the predecessor of LTE, leads to an averaging

of channel conditions and interference situations over the whole bandwidth. The use of OFDMA² in LTE creates a higher frequency selectivity of fading and interference. While the higher SINR variations make link adaptation more challenging, they also offer higher multi-user diversity gains if fast-enough feedback and channel adaptive scheduling is possible at the cost of a higher signaling overhead. Along the same lines, it has also been argued in the literature [JL09; LJ10] that in the context of modern RRM schemes in cellular systems like LTE, MIMO antenna configurations should better be used for spatial multiplexing rather than for transmit diversity: For low velocities, channel-adaptive scheduling and link adaptation are superior, and for high velocities, there is already an abundance of time and frequency selectivity so that additional diversity is superfluous [JL09; LJ10]. Nevertheless, transmit diversity offers a degree of resilience in situations that can occur in practice but are not modeled by system-level simulations: Individual Tx antenna chains at the BS can fail or situations where the propagation path between one of the Tx antennas and the MS is blocked (e.g., at the corner of buildings) can occur.

Summarizing our observations, we can conclude that SFBC is not a useful transmission mode to obtain high throughput when channel-adaptive transmissions are possible. It remains a viable backup solution especially in situations where no channel state information is available and only small amounts of data have to be transmitted.

5.4. Beamforming at the BS (MISO)

Before we evaluate the performance of MIMO with spatial multiplexing, we first discuss the performance of codebook-based beamforming in the downlink, also called rank-1 precoding in LTE. Obviously, multiple antennas at the BS are a prerequisite and we will use the antenna array layouts defined in Section 3.2.2 in the following. At the MS side, a single antenna is sufficient to receive a beamformed rank-1 transmission (MISO), but we will also evaluate the performance impact of adding more Rx antennas. As outlined in Chapter 2, beamforming is the multi-antenna technique of choice when spatial multiplexing is not feasible because the link has low SINR or a low-rank channel (e.g., due to LoS propagation and/or closely-spaced and thus correlated antennas). But as we will see, beamforming can be beneficial in a wide range of scenarios.

We start with a characterization of codebook-based beamforming in comparison to other beamforming or precoding schemes. After evaluating the potential gains under different array layouts and propagation conditions, we will show that beamforming operation poses additional link-adaptation challenges, which motivate the improvements that will be presented in Chapter 6.

²Note that in LTE there also exists the possibility to spread PRBs out in the frequency domain using so-called *virtual resource blocks*. We do not consider such schemes in the scope of this thesis.

5.4.1. Codebook-Based Beamforming in LTE (TM6)

When we discuss beamforming in this chapter, we have a MISO (see Fig. 2.1) channel in mind and consider precoding operations at the BS with a precoder $\mathbf{p} \in \mathbb{C}^{N_{\text{Tx}}}$ that maps a single layer of transmit symbols s onto the N_{Tx} transmit antennas. The precoder \mathbf{p} is also called a beamforming or weight vector [God97], whose complex weights are supposed to align the phases (and amplitudes) of the signal transmitted simultaneously from multiple Tx antennas in a way that they constructively interfere at the desired receiver. In order to accomplish this beamforming gain, the BS needs channel knowledge.

If the BS l' had perfect channel knowledge, e.g., in the form of the downlink MIMO channel matrix $\mathbf{H}_{l'k'}$ of user k' (or of the MISO channel vector $\mathbf{h}_{l'k'}^{\text{T}}$), it could use an *adaptive beamforming* technique that computes a tailored beamformer $\mathbf{p}_{k'}$ for the transmission to user k' . If it knew the channels to co-scheduled users (e.g., in other cells) the BS could even aim at selecting a precoder $\mathbf{p}_{k'}$ that minimizes the interference caused to these other users (*zero-forcing*). As outlined in Section 2.2.3, in transmission modes TM7 and TM8–10 LTE does support the use of arbitrary beamformers (or spatial multiplexing precoders in the latter cases) by allowing the use of user-specific demodulation reference symbols (DM-RS), which are beamformed/precoded in the same way as the PDSCH carrying the payload. However, perfect DL channel knowledge at the BS (CSIT) is not feasible in practice so that the BS has to rely on the codebook-based PMI feedback and/or on downlink–uplink channel reciprocity. The latter would be most promising in a TDD deployment using the same frequency band for downlink and uplink which we do not further consider here. Another approach that only requires CQI feedback is *opportunistic beamforming* [VTL02]: Beamformers are randomly applied and, based on SINR (CQI) feedback, those users who happen to benefit most from the random beam are scheduled. The idea is to further increase the multi-user diversity in order to realize even more multi-user diversity gain compared to the multi-user diversity from independently fading channels, which we analyzed in Section 5.1.1. One potential problem with opportunistic beamforming in a multi-cell system would be the randomly fluctuating inter-cell interference which could likely cause link adaptation difficulties, similar to those discussed at the end of this section.

The other class of beamforming in LTE, which is implemented in TM6 (and TM4), makes use of the LTE precoding codebook not only for feedback purposes but also utilizes the beamforming vectors defined in the codebook as beamformers for the downlink transmissions. This kind of beamforming is thus called *codebook-based beamforming*. It is also called *rank-1 precoding* because it uses the first column of a precoder \mathbf{P}_{PMI} as the beamforming vector $\mathbf{p}_{\text{PMI}} = \mathbf{P}_{\text{PMI}}[1]$, where the bracket notation means that the first column from matrix \mathbf{P}_{PMI} is selected. The use of rank-1 precoding is not restricted to MISO channels and it can, of course, also be used in MIMO channels whose rank is larger than one. But the effective channel $\mathbf{H}\mathbf{p}$ always has rank one (as long as $\mathbf{H} \neq \mathbf{0}$) regardless of the link being a MISO or MIMO link or the rank of \mathbf{H} .

The codebook-based beamforming in TM6 shares the same feedback design philosophy (see also the *Feedback* subsection in Section 2.2.3) as the closed-loop spatial multi-

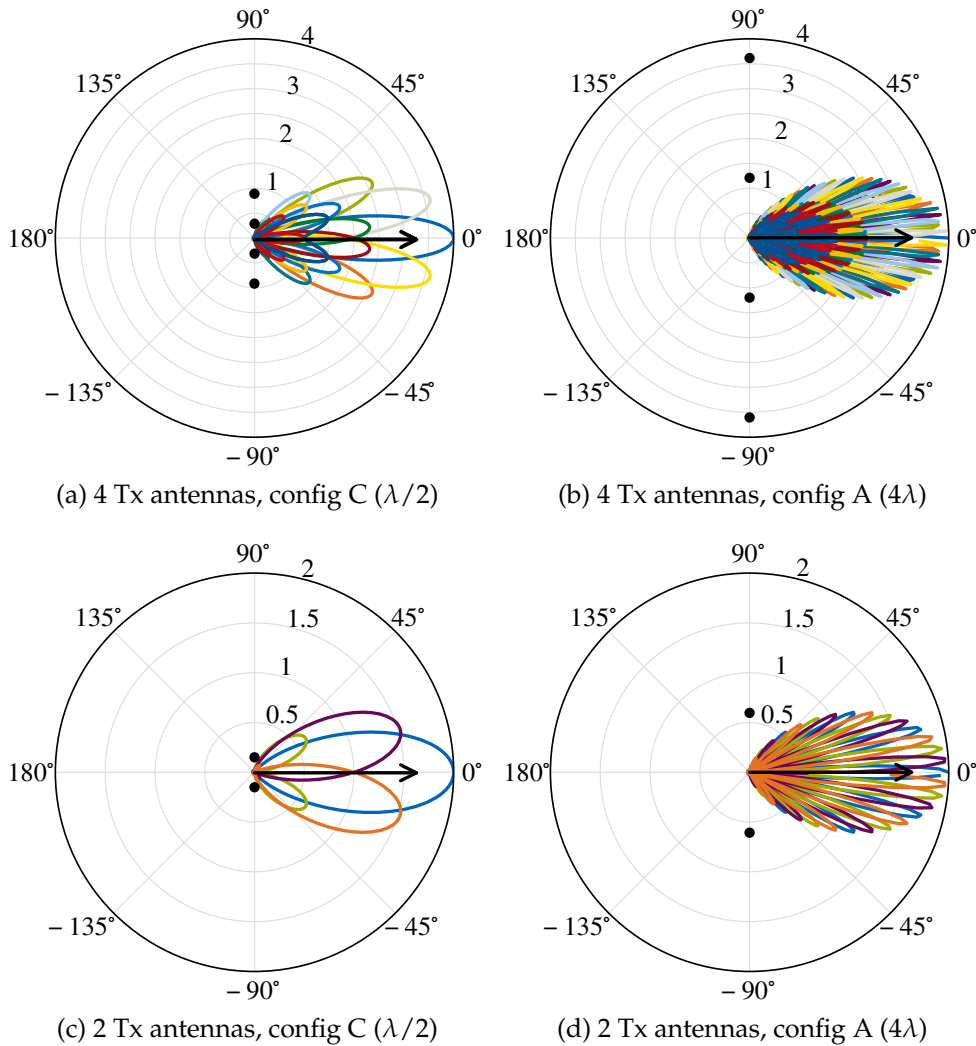


Figure 5.15.: Polar plots in linear scale of the rank-1 beamforming patterns from the LTE PMI codebook (see Tables B.3 to B.6 and Table B.7 in the Appendix) after application of the IMT-A azimuth sector pattern (see (3.3) and Fig. 3.3) with a boresight direction of 0° . Antennas are symbolized by black dots; see Fig. 5.16 for the PMI color code.

plexing in TM4, which is based on the same codebook: The MS estimates the pure DL channel matrix \mathbf{H} and determines which precoding vector $\mathbf{p}_{\text{PMI}^*}$ out of the possible beamformers from the codebook would yield the best effective channel $\mathbf{H}\mathbf{p}_{\text{PMI}^*}$. That choice is then signaled in the form of a PMI report to the BS. The BS then usually follows that recommendation without having further knowledge of the downlink channel \mathbf{H} . Note that TM6 is identical to TM4 but limited to rank-1 precoders. It only exists as a separate transmission mode because for rank-1 only operation, some feedback overhead (e.g., for the rank indicator) can be saved.

Figure 5.15 illustrates the beamforming patterns corresponding to the beamformers that are available in the LTE codebook. Given a beamforming vector, also the BS an-

tenna type and array layout influence into which directions the transmit power is emitted from the BS. In this figure, we only consider purely vertically polarized arrays, where the antennas are spaced at a distance of $\lambda/2$ and 4λ , corresponding to 3GPP array layouts C and A, respectively (see Fig. 3.5 on page 65). Note that with an antenna spacing of $\lambda/2$ (the two figures on the left), each codebook beamformer results in one well-defined beam of radiated energy. With a spacing of 4λ , each beamformer results in a beam pattern with multiple and more narrow peaks as shown on the right-hand side of Fig. 5.15.

A BS with two Tx antennas can only use four different beamformers whose patterns are plotted in Figs. 5.15c and 5.15d using different colors. With four Tx antennas, 16 PMIs are available in the codebook as listed in Tables B.3 to B.6 in the Appendix. The rank-1 beamformer is always the first column of the precoder matrix, denoted by $\mathbf{p} = \mathbf{P}[1]$. In the Appendix, we visualize the beam patterns corresponding to all four layers. Basically, the precoder for each layer can be interpreted as a beamforming vector for that layer. There are a few notable differences between the beam patterns plotted in the Appendix and the one shown in Fig. 5.15. The patterns shown in the Appendix are generated assuming an array with $\lambda/2$ spacing and without taking an antenna sector pattern into account. In contrast, in Fig. 5.15 the IMT-A sector pattern is superimposed, which effectively removes the sidelobes as well as the mirror images towards the other side of the linear array, whose antennas are represented by black dots in the figure. Another notable difference is that, in the case of 4 Tx antennas, not all of the 16 PMIs in the codebook produce beam patterns that appear to be useful. Looking at Fig. 5.15a, we can only identify seven different PMIs (see Fig. 5.16 for the PMI color code) that appear useful for beamforming because for no azimuth angle in the considered sector pattern, any of the nine other PMIs produces a stronger gain. However, in a multipath propagation environment, there might be significant propagation paths that depart from the BS with wider angular spreads than covered with one of the beams depicted in Fig. 5.15. This is especially true for the narrow beams resulting from the 4λ antenna spacing with config A. In the following, we will evaluate the codebook-based beamforming in different setups.

In Fig. 5.16 we show the spatial distribution of PMIs over the RMa scenario obtained from a system-level simulation with four different setups represented by the individual figures. In the simulations, we distribute thousands of mobiles at regularly spaced positions in the scenario, which is made-up of seven three-sector sites. Each pixel's color stands for the dominant wideband PMI feedback computed by the MS located at that position. The PMI feedback mechanism implemented in `IMTAphy` picks the PMI that delivers the highest post-filtering rate given the current interference situations on the PRBs in question. In general, it depends on the chosen receive filter (e.g., MRC, IRC, or MMSE) which in turn might depend on the estimated IPN covariance matrix at feedback computation time. In Fig. 5.16 we only consider a single Rx antenna at the MS so that Rx filtering is not possible.

We choose the RMa scenario for the illustration in Fig. 5.16 because, according to Table 3.3, with 82% it has the highest percentage of serving links that exhibit LoS propagation conditions. In LoS conditions, we expect the large-scale geographic prop-




Figure 5.16.: Geographic distribution of PMIs in the RMA scenario (high LoS prob.)

erties (i.e., the LoS direction relative to the array's broadside direction) to dominate the codebook-based beamformer selection. In fact, we observe in the two figures with $\lambda/2$ spacing on the left in Fig. 5.16 that in each of the 21 cell sectors, the geographic distribution of best PMIs corresponds to the main lobe directions of the $\lambda/2$ beam patterns shown in Fig. 5.15. While the different beam directions are clearly visible, a certain noise-like percentage of users remains who report arbitrary PMIs because, for example, they experience multi-path propagation with no or only a weak LoS component. Figure A.4 in the Appendix contains a similar visualization of the geographic PMI in the UMa scenario, which has the lowest percentage (37%) of LoS serving links. There, the geographic areas illuminated by different beams can still be identified in the case of $\lambda/2$ spaced Tx antennas. With 4λ spacing, as shown on the right-hand side of Figs. 5.16 and A.4, the geographic distribution of PMIs appears almost random due to the much narrower beams with multiple peaks as shown in Figs. 5.15b and 5.15d.

Table 5.13.: Simulation parameters for Figs. 5.17 and A.5 and Tables 5.14 and 5.15

Parameter	Value
Scenario	rural macro (RMa, see Fig. 5.17) and urban macro (UMa, see Fig. A.5)
Number of drops	5
Number of MS	on average $N_{MS} = 10$ users per cell
BS antennas	1, 2, and 4 antennas with purely vertical polarization and $\lambda/2$ (config C) as well as 4λ (config A) spacing
MS antennas	1, 2, 3, and 4 antennas with purely vertical polarization and $\lambda/2$ spacing
Transmission mode	TM6
Receive filter	MRC and IRC
Link adaptation	with outer-loop link adaptation (OLLA)
Scheduler	proportional fair (PF)
Channel estimation	perfect channel matrix and IPN covariance matrix estimation
Feedback	subband CQI and wideband PMI
Overhead	CRS ports overhead according to number of Tx antennas

Tables A.1 and A.2 in the Appendix list the frequencies with which the individual PMIs in the 2 Tx and 4 Tx codebook are reported in the RMa, UMa, and UMi scenario with antenna configs A and C. For the setup with correlated antennas at the BS ($\lambda/2$ spacing, config C) the percentages in Tables A.1 and A.2 reflect the visual impression from Figs. 5.16 and A.4: For 4 Tx antennas, only PMIs 0 . . . 7 (without PMI 2) and for 2 Tx antennas PMIs 0, 2, and 3 are frequently chosen. Overall, PMIs 8 . . . 15 are not frequently chosen for rank-1 precoding with the lowest probability (3%) in the LoS-rich RMa scenario with correlated antennas to the highest probability (34%) in the UMi scenario with uncorrelated (4λ spacing). We will exploit this concentration on a few frequently selected PMIs for an improved beamforming scheme that we will present in Section 6.2.1.

5.4.2. Performance of Codebook-Based Beamforming

For the performance evaluation of codebook-based beamforming in LTE, we focus on the two scenarios with the highest (RMa) and lowest (UMa) percentage of serving links with line-of-sight propagation conditions. Table 5.13 summarizes the relevant simulation parameters for the system-level simulations discussed in the following. We compare the performance with codebook-based precoding for both the 2 Tx and 4 Tx antenna codebook. As a reference, we also include single Tx antenna results. While the rank-1 precoding works with a single Rx antenna, we also investigate what performance gain MRC and IRC receive filters offer when two, three, or four Rx antennas with $\lambda/2$ spacing are available.

Figure 5.17 shows CDFs for the post-scheduling and post-receiver-combining per-user average SINR as well as the normalized user throughput in the RMa scenario for four

Table 5.14.: TM6 system spectral efficiencies in bit/s/Hz depending on the receiver type, antenna configuration, and antenna spacing in the RMa scenario

Receiver / array type	MIMO link (Tx \times Rx) antenna configuration											
	1 \times 1	1 \times 2	1 \times 3	1 \times 4	2 \times 1	2 \times 2	2 \times 3	2 \times 4	4 \times 1	4 \times 2	4 \times 3	4 \times 4
MRC, config C	0.84	1.20	1.41	1.55	1.11	1.47	1.66	1.78	1.46	1.80	1.97	2.08
MRC, config A	0.84	1.20	1.41	1.55	1.06	1.39	1.56	1.68	1.35	1.66	1.82	1.92
IRC, config C	0.84	1.48	1.99	2.37	1.11	1.79	2.26	2.59	1.46	2.16	2.58	2.85
IRC, config A	0.84	1.48	1.99	2.37	1.06	1.69	2.15	2.47	1.35	1.98	2.39	2.67

different setups each (MRC vs. IRC and BS antenna config A vs. C). In the Appendix, equivalent results for the UMa scenario can be found in Fig. A.5. According to the legend, the single Tx antenna case is always denoted by a solid line, whereas the 2 Tx and 4 Tx cases are denoted by dashed and dotted lines, respectively. Different line colors stand for different numbers of Rx antennas. When comparing the figures, it is helpful to use the black and solid lines as a reference: with a single Rx antenna, there is no difference between IRC and MRC (odd and even figure rows) and with a single Tx antenna (left and right-hand side) there is no difference between BS antenna array configurations A and C.

Similar to the SIMO situation evaluated before, in Figs. 5.17 and A.5 we observe that with every additional antenna at either side of the link the SINR grows. As expected, the gain is higher for IRC, but also with MRC the SINR CDF is shifted by over 12 dB when comparing the 1 \times 1 and 4 \times 4 setup. Naturally, the higher SINRs translate into higher throughputs even though the CRS overhead for increasing numbers of Tx antennas reduces the throughput growth a bit. While each additional antenna is beneficial, additional Rx antennas are much more valuable than additional Tx antennas. Even with simple MRC combining, a 1 \times 4 configuration with 1 Tx and 4 Rx antennas offers more throughput gain compared to the 1 \times 1 situation than a 4 \times 1 configuration with 4 Tx and 1 Rx antennas.

Table 5.14 lists the average per-cell spectral efficiencies in the RMa scenario, the corresponding results for the UMa scenario can be found in the Appendix in Table A.3. In the RMa scenario with closely spaced BS antennas (config C), we observe a 74% spectral efficiency gain when going from one (1 \times 1) to four Tx antennas in a 4 \times 1 setup. However, with a single Tx antenna but four instead of one Rx antennas, we observe a 182% gain when IRC filtering is used in the RMa scenario. If we then increase the number of Tx antennas also to four (4 \times 4), only an incremental gain of 20% is realized. There are two main reasons for this discrepancy between the benefits of Tx and Rx antennas. First, accurate channel estimation for MRC or IRC combining at the MS is much easier than finding a suitable beamformer at the BS when only very limited and delayed feedback in the form of highly quantized PMIs is available. Second, to facilitate channel estimation, the BS has to increase the CRS overhead when using more than one antenna, which limits the potential gains. Going from one to four CRS ports for 4 Tx antennas, the overhead grows by 6 percentage points (see Table 2.3) leading to an accordingly smaller throughput gain from using more Tx antennas.

5. Performance of Fundamental RRM Schemes

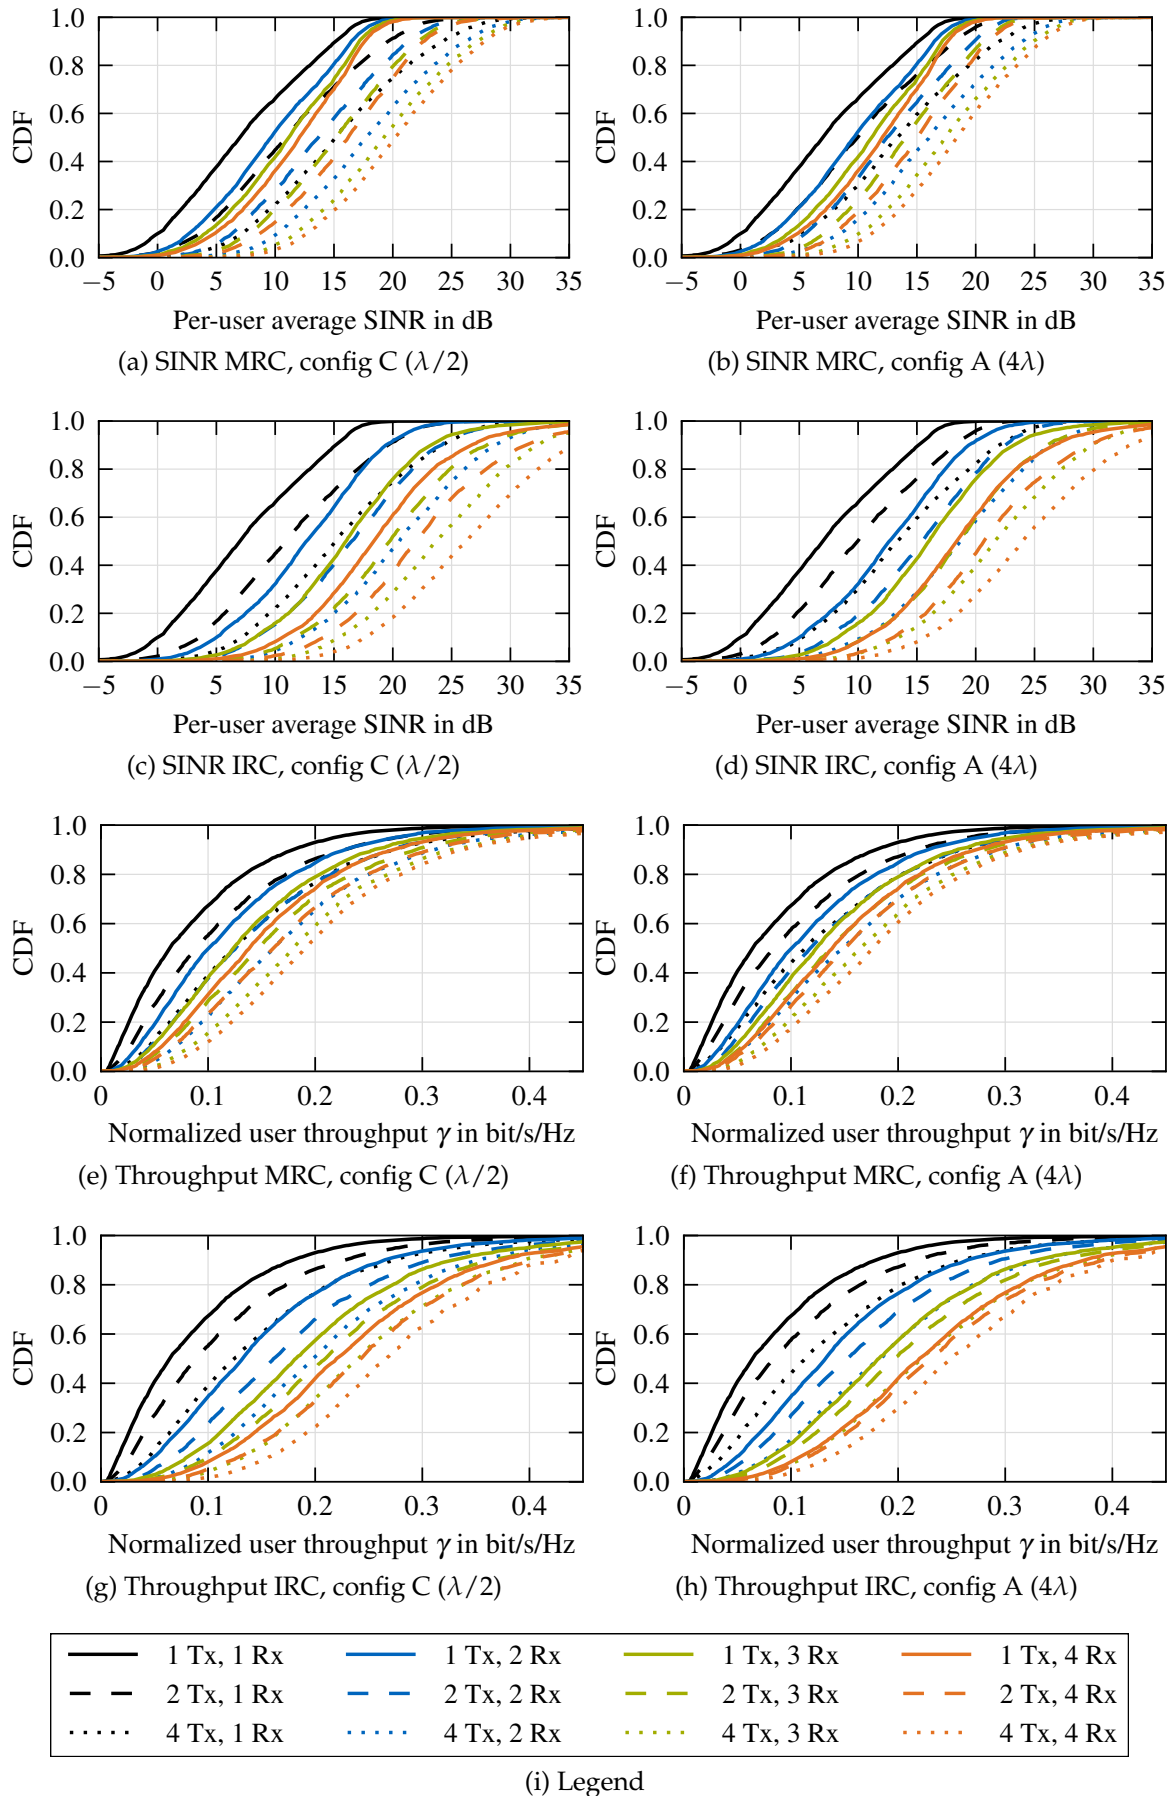


Figure 5.17.: Performance of TM6 in the RMa scenario

Table 5.15.: TM6 cell-edge user throughput in bit/s/Hz depending on the receiver type, antenna configuration, and antenna spacing in the RMa scenario

Receiver / array type	MIMO link (Tx \times Rx) antenna configuration											
	1 \times 1	1 \times 2	1 \times 3	1 \times 4	2 \times 1	2 \times 2	2 \times 3	2 \times 4	4 \times 1	4 \times 2	4 \times 3	4 \times 4
MRC, config C	0.011	0.025	0.034	0.041	0.019	0.035	0.044	0.049	0.031	0.053	0.065	0.074
MRC, config A	0.011	0.025	0.034	0.041	0.016	0.032	0.041	0.046	0.024	0.044	0.055	0.063
IRC, config C	0.011	0.035	0.062	0.084	0.019	0.051	0.080	0.101	0.031	0.074	0.106	0.130
IRC, config A	0.011	0.035	0.062	0.084	0.016	0.044	0.071	0.093	0.024	0.058	0.086	0.109

Also similar to the SIMO results discussed in Section 5.2, we observe an even higher throughput gain for the cell-edge users defined by the 5 %ile of the per-user throughput CDF. The cell-edge user throughputs in the RMa scenario are shown in Table 5.15 and results for the UMa scenario can be found in the Appendix. We observe that 4 Tx antennas in a 4 \times 1 link offer a two- to three-fold cell-edge user throughput gain compared to the 1 \times 1 setup. Together with spatial processing at the receiver, the gain can be almost 12-fold in the case of an IRC filter with correlated antennas in a 4 \times 4 setup compared to the 1 \times 1 case. Due to the logarithmic relationship between the channel capacity and the SINR (see (2.15) on page 15), SINR gains through beamforming at low SINR operating points (geometries) at the cell edge yield a proportionally larger throughput gain than the same SINR gains at higher geometries.

The SINR CDFs also exhibit another interesting effect, which we already observed in Section 4.5.3. The additional SINR gain with the MRC filter is much smaller at high SINRs than for the rest of the SINR range. Since high SINR users are often found close to the BS, their SINR is limited by the interference coming from the backside of two co-located sector antennas (with an attenuation of -20 dB, see Section 3.2.1). In that situation, the MRC filter that is tailored to the (often LoS) serving channel is not very well suited because it also amplifies these interferers resulting in almost no gain. The IRC filter handles this situation better since it is aware of the interference and aims at attenuating the interferers, for example, by not focusing on the LoS ray but rather on weaker multipaths from the co-located sectors exhibiting uncorrelated interference.

What is important to note is that while correlated antennas at the BS (config C, $\lambda/2$ spacing) produce beams with a clearly visible directivity, the seemingly erratic beam distribution with the uncorrelated antennas (config A, 4λ) produces qualitatively identical gains, which are only slightly lower. A similar observation can be made when comparing the LoS-rich RMa scenario with the NLoS-rich UMa scenario. The former benefits only slightly more from the beamforming than the latter.

The PMI feedback in IMT_{Aphy} is selected by picking that precoding vector $\mathbf{p}_{\text{PMI}^*}$ whose average effective channel $\mathbf{H}\mathbf{p}_{\text{PMI}^*}$ promises the best post-receiver rate when averaged over all PRBs considered for the feedback computation. Here, we only considered wideband PMI feedback. In LTE, and also in IMT_{Aphy}, it is possible to signal PMI subband feedback, where the best precoder per subband is determined based on the average over the PRBs of that subband only. We do not show results for subband PMI here because the benefit from subband-selective PMI was very limited in our eval-




Figure 5.18.: Illustration of the flashlight effect

uations with the feedback algorithm implemented in `IMTAphy`. Similar observations were published in the literature [BA09].

To summarize this codebook-based performance evaluation section, we can conclude that both correlated and uncorrelated antenna array layouts at the BS are suitable for beamforming in TM4 and TM6 even though correlated antennas yield slightly larger gains. The same is true for LoS-rich scenarios like RMA, but also in the UMA scenario we see substantial gains compared to single antenna transmission. The scenarios we considered so far were dense (i.e., with wraparound) outdoor cellular scenarios characterized by relatively low SINRs (see also the geometry, i.e., long-term 1×1 SINR distribution in the IMT-A scenarios in Fig. 3.8b on page 77). In other scenarios that are characterized by higher geometries, such as the InH scenario, we expect the beamforming throughput gains to be much more limited because for SINRs well above 20 dB we have reached the highest MCS in LTE where additional SINR gains cannot improve the single layer throughput. Under those conditions, we will have to resort to spatial multiplexing as discussed in Section 5.5.

5.4.3. Link Adaptation for Beamforming in Multi-Cell Systems

Before we conclude the discussion of beamforming in LTE, we want to highlight one challenging system-level aspect. When dynamically applying beamformers for individual users, the power emitted by a BS into a certain solid angle can change with the scheduling decisions from PRB to PRB and from TTI to TTI. Compared to a non-beamforming system, where the antenna pattern gain for a certain combination of azimuth and elevation angles is fixed, this so-called *flashlight effect* leads to highly fluctuating levels of inter-cell interference in neighboring cells. The situation is depicted in Fig. 5.18: A mobile is served by its base station using a suitable serving beam, colored blue in this example. A neighboring base station (on the right) also uses beamforming to serve its users (not depicted here) and can choose between multiple beams, among them the two beams shown in Fig. 5.18. The associated problem arises when the beamformed interference situation at the time of the CQI feedback computation is different from the beamformed interference during the transmission whose scheduling and link

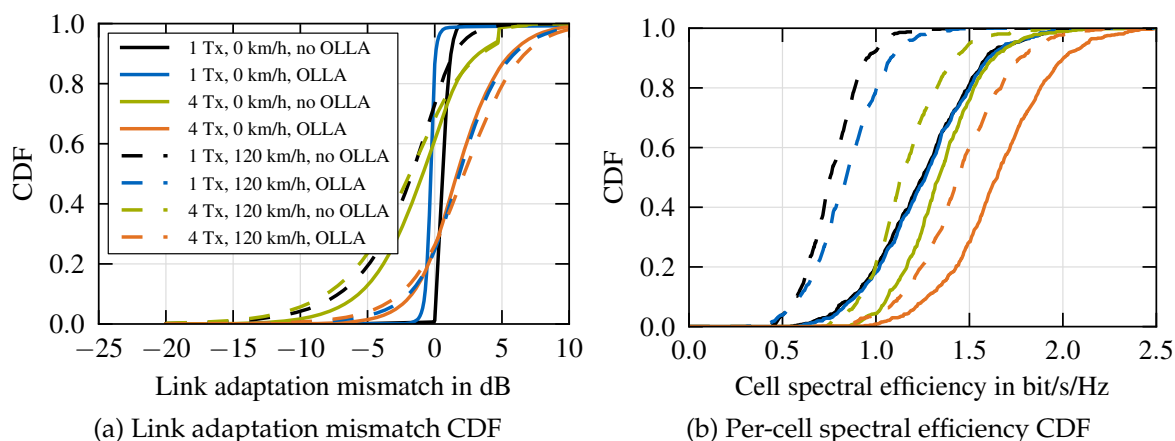


Figure 5.19.: Link adaptation mismatch and per-cell system spectral efficiency distributions demonstrating the flashlight effect in the RMA scenario

adaptation is based on that CQI feedback. Two cases can be distinguished: If during feedback computation on the considered PRBs the MS is not interfered by the neighboring cell because that BS uses the non-colliding green beam for transmission to some user, the MS will report a certain CQI to its serving BS. Some TTIs later, when the BS has processed the CQI and has scheduled a transmission to the MS, the neighboring BS might use the orange beam to serve some user. Since the interference level is now much higher due to the strong interfering beam, the transmission based on the old optimistic CQI is likely to fail. In the other case, the MS is interfered by the strong beam during CQI computation and reports a correspondingly lower CQI to its base station. If, during the transmission with an accordingly lower MCS, the neighboring cell then uses the non-colliding beam, the transmission is likely successful, but probably does not exploit the potential rate because it is based on a too pessimistic CQI.

This effect can be demonstrated in the link adaptation mismatch plot, which we introduced in Section 5.1.2. Figure 5.19 shows the link adaptation mismatches in the RMA scenario, where we use a single Rx antenna and perform simulations (similar to those in the previous subsection) with codebook-based beamforming using four antennas and a single antenna as a reference. The results with and without outer-loop link adaptation are shown. In order to isolate the mismatch due to the flashlight effect from the fast fading influence of the channel, we conduct simulations with the RMA default speed of 120 km/h (dashed lines) and with (unrealistic) static channels (solid lines). In the static scenario with a single Tx antenna, Fig. 5.19a shows very limited link adaptation mismatches which can be reduced further by OLLA. When using beamforming with 4 Tx antennas and even with static channels, almost 80% of the transmissions scheduled by the greedy PF scheduler overestimate the SINR by up to 10 dB while the remaining 20% underestimate the SINR (here, the mismatch is capped at 5 dB by the effective SINR model). We see that the link adaptation mismatch from the fast fading in the RMA scenario leads to a slightly larger overestimation but to less underestimation (the fading rarely improves further when trying to schedule onto peaks). On the

right-hand side, Fig. 5.19b displays the distribution of per-cell spectral efficiencies. In all situations with a high SINR fluctuation (either due to fading or the flashlight effect) the OLLA mechanism is able to improve the throughput by reducing the SINR overestimation though at the expense of more underestimation. Once again, we observe a significant throughput gain with beamforming, even if OLLA is not enabled.

An obvious improvement to uncoordinated beamforming in multi-cell systems is to coordinate the beamforming between neighboring cells to avoid the flashlight effect. There have been numerous proposals in the literature and basic support has been added to LTE Rel-11 in connection with TM10. Besides the signaling necessary to coordinate between separate sites, also the additional channel knowledge (of the interfering channel) is a crucial factor for any coordinated beamforming scheme. In Section 6.2.1 we will present a simple implicit coordination scheme that already achieves substantial gains by avoiding the flashlight effect.

5.5. Spatial Multiplexing (MIMO)

In this section, we will now consider spatial multiplexing in LTE, that is, using multiple antennas at the BS and MS side (MIMO) for the transmission of multiple layers per user. We present it after our discussion of SIMO (Section 5.2) and MISO (Section 5.4) because it can be considered as a combination of these two techniques: The transmission of more than one layer via multiple Tx antennas in LTE makes use of pre-defined precoding matrices P at the BS where each column maps the symbols of its layer to the available Tx antennas. At the receiver-side, we rely on two linear MMSE receiver filter variants that were introduced in Section 4.2.2. Their naming corresponds to the MRC and IRC filters we considered for the reception of a single layer in Section 5.2: The MMSE-MRC filter is able to receive multi-layer transmissions by using multiple Rx antennas to suppress inter-layer interference but does not consider inter-cell interference. The MMSE-IRC filter, on the other hand, in addition to suppressing inter-layer interference, also allows to suppress inter-cell interference by taking an estimate of the full IPN covariance matrix into account.

The differences between precoding with open-loop operation (TM3) and closed-loop operation (TM4) will be briefly described in the next section. In Section 5.5.2 we will then present a performance evaluation where we distinguish between the MMSE-MRC and MMSE-IRC receivers. In addition, we will examine the impact of different propagation conditions in the IMT-A scenarios and evaluate the influence that antenna array design choices have.

5.5.1. Open-Loop and Closed-Loop Spatial Multiplexing in LTE

Closed-Loop Spatial Multiplexing in LTE (TM4)

The basic mode of operation of closed-loop spatial multiplexing in TM4 is identical to codebook-based rank-1 precoding in TM6 (and TM4). The MS returns the PMI of the

desired codebook entry. In contrast to Section 5.4, now that PMI refers to a precoding matrix where the number of columns corresponds to the desired number of layers. What is new compared to rank-1 precoding, is that the desired rank (i.e., the desired number of spatial layers) is additionally signaled by the rank indicator (RI) feedback (see also Section 2.2.3 for an introduction). In the 4 Tx codebook there are 64 different combinations of PMIs and ranks from 1–4 available (see Tables B.3 to B.6), while in the 2 Tx codebook (see Table B.7) only six different combinations exist.

Similar to the receiver design, the PMI and RI feedback computation strategies are not standardized and are vendor-specific. In IMT_{Aphy} , the PMI and RI feedback is computed by an exhaustive search. Given the serving channel estimate and interference situation at the time of feedback computation, all effective channels $\hat{H}P_{\text{PMI}}$ are fed into the employed linear receiver (i.e., either the MMSE-MRC or the MMSE-IRC filter) and the resulting per-layer SINRs are computed. The SINRs per subband are compressed using a link-to-system model and converted into an expected spectral efficiency, taking the transport block to layer mapping and the upper limit due to the highest MCS efficiency into account. The whole procedure is performed by each MS for all PRBs and subcarriers. For the feedback, the frequency granularity has to be considered as well. The RI can only be signaled wideband with a single RI report for the whole bandwidth. The PMI can either be configured wideband as well or, alternatively, it can be signaled frequency-selective where multiple PRBs form a subband for which one PMI value is reported. In IMT_{Aphy} , we thus pick the rank that gives the best throughput over the whole bandwidth considering the best PMIs per subband for that rank. The PMI feedback, which is updated more often, is then restricted to the current rank.

Considering the 4 Tx codebook in Tables B.3 to B.6, we see that for all ranks smaller than four, a subset of columns from the full matrix $P[1234]$ is selected. Note that, depending on the PMI, the order of columns might be permuted compared to the $P[1234]$ matrix. Similar to the rank-1 precoding case, we can associate each layer with a corresponding radiation pattern that depends on the antenna array layout. The patterns resulting for a linear array with $\lambda/2$ spacing are printed in Tables B.3 to B.6. Here, we want to highlight a special property of the LTE codebook. The columns within a precoder matrix are orthogonal and the full precoding matrices are (up to a constant factor) *unitary matrices*. By definition, their covariance matrix is thus the (scaled) identity:

$$P[1234]P[1234]^H = \frac{1}{4}I_4. \quad (5.7)$$

The total radiation pattern of a rank-4 transmission using all four layers is thus omnidirectional (see the black unit circle representing the sum pattern in the codebook tables in the Appendix) and only depends on the BS sector pattern. It thus prevents the flashlight effect observed with rank-1 precoding and creates a stable system-wide interference situation that is beneficial for link-adaptation in neighboring cells. For transmissions with less layers, the sum patterns do exhibit a PMI-specific directivity so that again a flashlight effect can happen. The highest directivity is obtained with a single layer transmission in TM4, which is exactly the rank-1 precoding case discussed before in Section 5.4.

Using precoding with closed-loop spatial multiplexing in TM4, the BS can achieve a Tx-side gain by focusing the energy towards the eigenmodes of the channel, as far as the coarse channel quantization provided by the PMI feedback allows that. It is, however, not necessary to use a specific precoder for spatial multiplexing transmission. In fact, the additional gain of CSIT and a corresponding spatial processing at the transmitter is rather small and vanishes at high SINRs (see Section 2.1.2 or the literature [PNG03] for a more detailed treatment). In the absence of CSIT, the identity matrix could be used as a “precoder”, but in the open-loop spatial multiplexing implemented in TM3 in LTE, a combination of alternating precoders is used to increase the diversity.

Open-Loop Spatial Multiplexing in LTE (TM3)

The open-loop spatial multiplexing transmission mode three (TM3) has been designed for situations where channel conditions (i.e., SINR and rank) support spatial multiplexing, but no closed-loop PMI feedback is provided, for example, due to high user mobility. As we have seen in Section 5.1.2, frequency-selective scheduling does not promise gains if CQIs are outdated due to higher speeds. Thus, TM3 is not designed for exploiting scheduling gains, but rather is made robust by providing different kinds of transmit diversity. If the MS signals rank-1 channel conditions, the BS employs the SFBC transmit diversity scheme as described in Section 5.3 for TM2. For a spatial multiplexing transmission, that is, for ranks ≥ 2 , an alternating set of precoders is used between adjacent resource elements, transforming spatial diversity into frequency diversity [BA09]. For some REs, the assigned precoder will be beneficial, while for others it will not fit well to the channel. Since even the smallest transport (code) blocks are encoded over a large number of REs, an average effective SINR value per transport block results. This averaging effect protects the transport blocks against severe fading or a beamformer mismatch. At the same time, it also precludes significant channel-adaptive scheduling gains because channel variations are averaged out.

In TM3, the precoders $\mathbf{P}_{M, N_{\text{Tx}}}(i)$ for a certain data symbol i (note that symbols are allocated in a frequency-first manner onto the available REs, see Fig. 2.5 on page 26) are computed using the following relationship depending on the number of transmitted layers M and the number of Tx antennas [36.211]:

$$\mathbf{P}_{M, N_{\text{Tx}}}(i) = \mathbf{W}_{N_{\text{Tx}}}(i)\mathbf{D}_M(i)\mathbf{U}_M, \quad (5.8)$$

where the DFT matrix \mathbf{U}_M and the matrix factor $\mathbf{D}_M(i)$ are defined as listed in Table B.2 in the Appendix. The precoder $\mathbf{W}_{N_{\text{Tx}}}(i)$ is the scaled identity matrix with $\mathbf{W}_2(i) = 1/\sqrt{2}\mathbf{I}_2$ if 2 Tx antennas are used. For 4 Tx antennas, $\mathbf{W}_4(i) = \mathbf{P}_k$ and one of four precoders \mathbf{P}_k from the 4 Tx codebook is chosen where the codebook index k for Table B.6 is computed as $k = \lfloor \frac{i}{M} \rfloor \bmod 4 + 12$.

The multiplication of $\mathbf{D}_2(i)\mathbf{U}_2$ yields $(1 \ 1)^T$ and $(1 \ -1)^T$ as the per-layer precoding matrix columns (beamformers) for 2 Tx antennas, where the layer-to-beamformer mapping alternates between each subcarrier. The beamformers are identical to the precoder columns of PMI 1 in the 2 Tx codebook, whose beam patterns are plotted in

Table 5.16.: Simulation parameters for the simulations presented in Section 5.5.2

Parameter	Value
Scenario	RMa and InH, results for UMa and UMi shown in the Appendix
Number of drops	10 drops for hexagonal scenarios, 75 drops for InH
Number of MS	on average $N_{MS} = 10$ users per cell with scenario-default speed
Subcarrier modeling	12 subcarriers per PRB modeled, both for TM3 and TM4
BS and MS antennas	four BS and MS antennas each; four different BS–MS antenna array configuration combinations: C–C, A–C, E–E, D–E (see Fig. 3.5 on page 65 for a reference)
Receiver type	MMSE-MRC (diagonal IPN covariance matrix) and MMSE-IRC (perfect IPN covariance matrix)
Transmission modes	TM1, TM3, and TM4
Link adaptation	with outer-loop link adaptation (OLLA)
CQI subband size	subband (i.e., 3 PRBs per subband)
PMI subband size	wideband (i.e., 50 PRBs per subband)
Scheduler	proportional fair
Overhead	4 CRS ports for TM3 and TM4, 1 CRS port for TM1

Table B.7. Similarly for 4 Tx antennas, the product of $\mathbf{D}_M(i)\mathbf{U}_M$ leads to a uniformly rotating distribution of the M layers to M beamformers. In addition, M subsequent groups, consisting of M subcarriers each, are precoded with the alternating precoders $\mathbf{W}_4(i)$ to add another degree of SINR variance across all REs in a code block.

To summarize, we highlight that TM3 allows spatial multiplexing without the need for closed-loop PMI feedback. It achieves an SINR averaging effect by rotating deterministically through a pre-defined set of precoders. That way, it is especially suited for high mobility contexts or if only a limited feedback overhead is desired. In addition to a single CQI report (both transport blocks are expected to find the same mix of SINRs), the rank needs to be determined, which is implemented in IMT_{Aphy} identically to TM4 with the difference that fixed precoders are applied per subcarrier according to the described scheme. Similar to the TM2 modeling in IMT_{Aphy} , a TM3 simulation requires modeling the channel with subcarrier resolution (i.e., 12 subcarriers per PRB with a 12-fold amount of channel data generated).

5.5.2. Performance of Spatial Multiplexing with TM3 and TM4 in LTE

For the performance evaluation of spatial multiplexing in this subsection, we focus on MIMO configurations with four transmit and receive antennas. We distinguish between the MMSE-MRC and the inter-cell interference suppressing MMSE-IRC linear receivers. In our simulations, the number of Rx antennas at the MS is larger than what is going to be deployed in the field in the foreseeable future. But with fewer antennas, we could not achieve spatial multiplexing transmissions with the maximum number of spatial layers supported in LTE Rel-8.

Table 5.16 lists the simulation assumptions. We consider all IMT-A scenarios and compare the results obtained with different antenna configurations. Both the propagation environment and the antenna configuration are expected to influence the channel correlation and thus the suitable number of layers and ultimately the throughput performance. The antenna configurations at the BS are again denoted by the 3GPP-defined [36.814] letters, which are illustrated in Fig. 3.5 on page 65. Configurations C and A use purely vertically polarized antennas, whereas E and D use cross-polarized antennas. C and E have an antenna spacing of $\lambda/2$ and A and D use 4λ . On the MS side, we use the same nomenclature to distinguish the polarization setup, and we always stick to a $\lambda/2$ spacing.

For the comparison of spatial multiplexing performance and the discussion of the rank statistics, we distinguish between the performance in outdoor interference-limited hexagonal cell deployments (i.e., the UMa, UMi, and RMa IMT-A scenarios) and the performance in indoor hotspot environments (i.e., the InH scenario). We choose the RMa scenario as the representative for the outdoor deployments because it is the scenario with the lowest percentage of higher ranks, which is in stark contrast to the InH scenario. Table A.5 in the Appendix summarizes the system spectral efficiency and cell-edge user throughput performance of the RMa and InH scenarios discussed in the following. The corresponding results for the UMa and UMi scenarios can be found in the Appendix in Table A.6 and Figs. A.6 and A.7.

Spatial Multiplexing in Outdoor Interference-Limited Scenarios

The upper part of Fig. 5.20 shows the percentage of transmission ranks 1, 2, 3, and 4 as chosen by the IMT_{Aphy} feedback routine averaged over all users and all feedback computation instances. In the lower part, the corresponding user throughput distributions with open-loop (TM3) and closed-loop (TM4) spatial multiplexing are shown. From Figs. 5.20a and 5.20b, we observe that the antenna array configurations at the BS and MS side heavily influence the preferred transmission ranks. The influence from the antenna spacing has a smaller impact compared to the use of purely vertically polarized versus cross-polarized antennas. The latter almost always allow a minimum of rank-2 due to the small coupling and hence uncorrelated fading between orthogonal polarization directions. However, overall the channel conditions in the IMT-A outdoor scenarios only rarely allow the maximum of four layers that could be supported by the 4×4 antenna setup. Smaller differences are observed between TM3 and TM4. In general, the straight-forward rank selection implemented in IMT_{Aphy} seems to prefer slightly lower ranks with TM4 than with TM3 in the RMa scenario.

Bigger differences between TM3 and TM4 are observed in the resulting throughputs. For all antenna configurations, the closed-loop spatial multiplexing with TM4 provides higher throughputs than the open-loop TM3 operation in the outdoor scenarios. However, in comparison with the respective TM1 throughput, for which a slightly lower overhead due to only a single CRS port is modeled (6% delta, see Table 2.3), TM3 and TM4 offer only small to medium gains. In case of TM3 operation, especially with co-polarized antennas at the BS or interference rejection combining MMSE-IRC receivers,



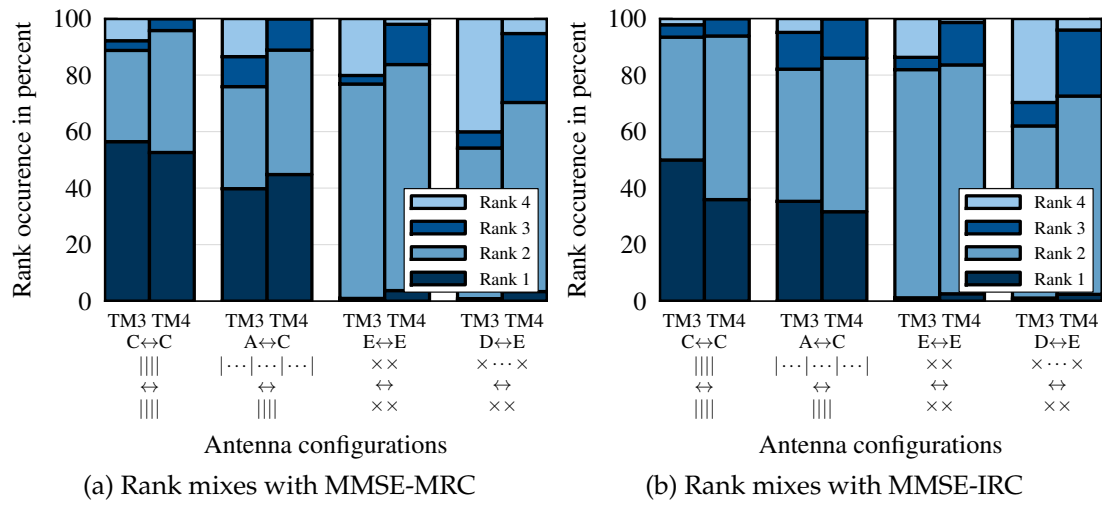


Figure 5.20.: Impact of 4×4 MIMO antenna configurations on the rank mixes and per-user throughput distribution in the RMa scenario; four antenna array combinations (BS \leftrightarrow MS, see Fig. 3.5 on page 65 as a reference), each for TM3 and TM4 as well as TM1 as a reference for the throughput

the performance is even worse than the simple single-antenna transmission with TM1. There are four main reasons for the poor performance of spatial multiplexing in outdoor scenarios:

First, the majority of users in the outdoor scenarios RMa, UMa, and UMi experience only low to medium SINRs with about 70% of users having an SINR of less than 10 dB as shown in Fig. 3.8b on page 77. Since the Tx power has to be distributed equally over the layers, and additionally inter-layer interference occurs, the spatial-multiplexing SINRs and thus the feasible rates per layer can be much lower than for single-layer transmissions. The available per-layer precoders usually also do not perfectly align with the eigenmodes of the channel so that low per-layer SINRs can result. Based on the waterfilling argument, additional layers should only be used after allocating power to the strongest layer (in terms of SINR). Consequently, the feasible sum throughput of a multi-layer transmission only surpasses the possible single-layer throughput when all layers can significantly contribute to the sum throughput, which is more likely at high SINRs. For example, Holma [HT11] claims that in a 2×2 LTE system, spatial multiplexing outperforms single-layer transmissions only for SINRs larger than 12 dB. Accordingly, we observe that spatial multiplexing throughput gains are not achieved for all users, but especially the better-served users benefit as seen for the higher performance bins of the throughput distribution in Figs. 5.20c and 5.20d.

Second, the outdoor scenarios exhibit a significant amount of LoS propagation conditions on the serving links. LoS makes the channel more correlated and thus lowers the feasible ranks and the MIMO capacity. Cross-polarized antennas help to reduce the correlation, which explains the better performance of antenna configurations E and D as shown in Figs. 5.20c and 5.20d. Third, in a simulation scenario where for all links in the multi-cell system the same transmission mode is applied, transmissions modes 3 and 4, with a certain percentage of spatial multiplexing transmissions in the system, lead to an interference situation where more independent interfering layers are present. This makes it harder for the MMSE-IRC receiver to suppress interference as easily as in the TM1 case, where each interferer only transmits a single layer. For that reason, TM1 is even more competitive when the MMSE-IRC receiver is used. Fourth, compared to TM1, interfering spatial multiplexing transmissions with less than full rank contribute to an interference flashlight effect that makes link adaptation more unreliable, which is detrimental to the spectral efficiency as discussed in Section 5.4.3.

Comparing the performance of rank-1 precoding with a 4×4 antenna configuration in the RMa scenario (as shown in Fig. 5.17 and Tables 5.14 and 5.15) to the TM4 performance listed in Table A.5 in the Appendix, we observe a slight system spectral efficiency advantage for TM4 compared to TM6. This is expected because the TM4 feedback algorithm in the mobiles can opt for a rank-1 transmission equivalent to TM6 so that the performance of TM4 should not be worse than that of TM6. However, the interference situation cannot be controlled by the UE feedback strategy so that, theoretically, the MS might report a rank ≥ 2 even though it would have preferred a system-wide limitation to TM1, which it cannot enforce on its own. Especially for cell-edge users equipped with IRC receivers, rank-1 interference is more advantageous as reflected by a better TM6 cell-edge user throughput compared to TM4 and TM3. We

also observe that at the cell-edge, correlated BS array setups (i.e., configuration C) offer the highest throughput because they provide the best beamforming gain allowing to boost cell-edge SINRs.

Overall, we have seen that in outdoor scenarios the potential gains from spatial multiplexing are limited and far below the theoretical limit of a gain factor of $\min(N_{Tx}, N_{Rx})$. For users with higher SINRs, the use of cross-polarized antennas proves beneficial because it lowers the correlation and increases the spatial multiplexing gain. For users with low SINR, however, co-polarized BS antennas are better so that the operator has to make a deployment decision in favor of one of the two groups of users.

Spatial Multiplexing in Indoor and Hotspot scenarios

The situation is very different in hotspot deployments as represented by the IMT-A indoor hotspot (InH) scenario. These deployments are not as interference-limited as the outdoor scenarios so that much higher SINRs are achieved. According to Fig. 3.8b on page 77, almost all users experience positive long-term SINRs and for more than 40 % of users the SINR exceeds 20 dB. Under these conditions, which are further improved by the 4 Rx antennas assumed here, the throughputs achieved by TM1 are limited by the maximum single-layer MCS efficiency of about 3.77 bit/s/Hz in LTE (maximum single-layer transport block size of 75 376 bits in a 20 MHz bandwidth). Figures 5.21c and 5.21d show this limitation of the TM1 throughput distribution and demonstrate that, under these conditions, spatial multiplexing with both TM3 and TM4 can achieve a substantial throughput gain by using a richer mix of spatial layers as illustrated by Figs. 5.21a and 5.21b. No big performance difference between closed-loop and open-loop operation is observed with a slight advantage for TM3. A closer inspection (not shown) reveals that two opposite effects balance each other: TM4 achieves higher SINRs than TM3 due to channel-adaptive precoding, but the inherent averaging effect of TM3 allows a better link adaptation due to a smaller link adaptation mismatch. We note that despite a high serving link LoS probability of 92 % and the corresponding higher channel correlations, the high SINRs in the InH scenario allow a rich mix of transmissions ranks. Again, we observe that cross-polarized antenna configurations perform better due to lower correlations and thus increased ranks.

We can summarize our performance evaluation of TM3 and TM4 like this: The performance of spatial multiplexing heavily depends on the deployment scenario and the associated propagation conditions and SINRs. For outdoor scenarios, spatial multiplexing is often not beneficial (in comparison to TM1) because channel conditions do not favor higher ranks. In this context, the antenna array configurations at the BS play an important role because with cross-polarized antennas, at least a rank-2 transmission can be achieved. Unlike other RRM efficiency determinants, the BS array setup is a long-time choice that cannot be adapted during runtime like the transmission mode or rank selection. In contrast to MS design choices (e.g., antenna configurations and the receiver type), it can be adapted to the deployment scenario though. However, the BS array choice in the field is often constrained by form factor limitations which fa-

5. Performance of Fundamental RRM Schemes

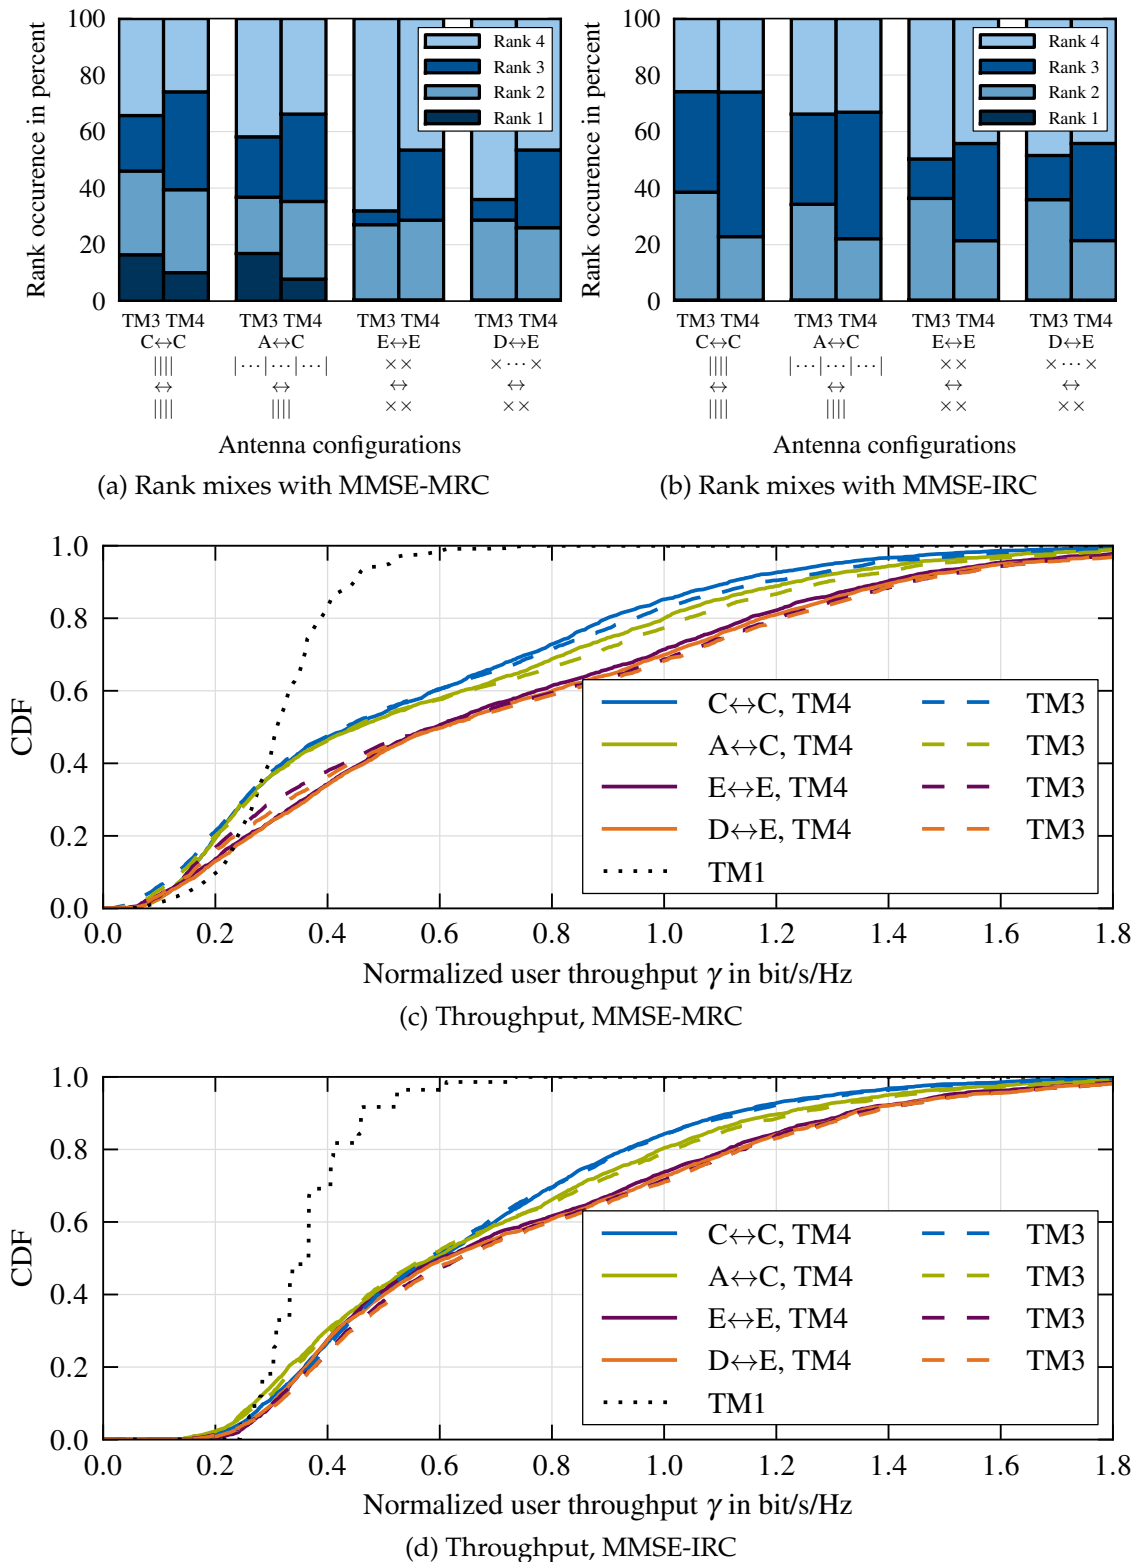


Figure 5.21.: Impact of 4×4 MIMO antenna configurations on the rank mixes and per-user throughput distribution in the InH scenario; four antenna array combinations (BS↔MS, see Fig. 3.5 on page 65 as a reference), each for TM3 and TM4 as well as TM1 as a reference for the throughput

vor compact setups so that the $\lambda/2$ -spaced cross-polarized setup E appears especially attractive given its performance results.

Simulation-Based PHY Overhead Evaluation at the Example of TM4

In Section 2.3, we discussed the different kind of losses occurring on the physical layer that prevent the link spectral efficiency from reaching the channel capacity. Based on general LTE overhead factors and link performance values published in the literature for a single SINR operating point, we computed the total PHY layer loss in (2.32) on page 46. Using TM4 as an example and with the help of `IMTAPHY` system-level simulations, we now want to evaluate the PHY losses across the whole range of link conditions (i.e., SINRs, MIMO channels, interference situations, etc.) in a typical IMT-A scenario.

We conduct system-level simulations with the parameters defined in Table 5.16. As outlined in Chapters 2 and 4, our simulator already takes all overhead aspects discussed in Section 2.3 into account when it determines the normalized user throughput $\gamma_{k'}$ for a user k' . In order to compute the total loss incurred, we determine the mutual information rate $\text{MI}_{t'k'}$ on the link from the transmitting BS t' to the MS k' . The equation was already introduced in (2.27) on page 23, but we repeat it here for convenience:

$$\text{MI}_{t'k'} = B \log_2 \det \left(\mathbf{I}_{N_{\text{Rx}}} + P_{\text{tx},k'} \mathbf{H}_{t'k'} \mathbf{P}_{k'} \mathbf{P}_{k'}^H \mathbf{H}_{t'k'}^H \mathbf{R}_{\text{IPN},k'}^{-1} \right) \text{ bit/s.} \quad (5.9)$$

From the throughput $\gamma_{k'}$ (normalized to the system bandwidth B) and from the $\text{MI}_{t'k'}$ per user k' , we compute the PHY loss for that user as $1 - \frac{B \cdot \gamma_{k'}}{\text{MI}_{t'k'}}$ and average over all users. Note that the loss definition here is not exactly identical to the one in Section 2.3 because the literature values in Table 2.2 on page 43 take the single-cell channel capacity as their benchmark. Here, we compute the mutual information based on the precoders \mathbf{P}_k that were chosen during the normal TM4 simulation. That means we neither model the codebook-based precoding loss in row b) of Table 2.2 nor include losses from missing waterfilling in the time, frequency, or spatial domain. Furthermore, no potential multi-cell optimizations yielding some kind of network capacity are taken into account since all interfering precoders are as chosen during the simulation where no form of inter-BS coordination is applied.

Table 5.17 lists the PHY losses obtained in the UMi and InH scenarios with different MIMO configurations and MMSE-MRC as well as MMSE-IRC receivers. Compared to the single result of 65.5% presented in (2.32), which was computed³ for a 4×4 MIMO link at 10.4 dB, the presented losses are in a similar range from 56.6% to 81.7%. The differences have several reasons: With a higher number of Tx antennas, the CRS overhead grows accordingly (see Table 2.3 on page 45). Also, the better-performing MMSE-IRC receivers have lower losses due to their better performance. We consider the scenario-specific mix of SINRs, which leads to a whole range of different operating

³the equivalent result for the 1×1 setup is $1 - (1 - 0.291) \cdot (1 - 0.256) \cdot (1 - 0.1) \cdot (1 - 0.067) = 55.7\%$, based on the values in Tables 2.2 and 2.3

Table 5.17.: PHY losses in the UMi and InH scenario determined in system-level simulations as $1 - \frac{B \cdot \gamma}{\text{MI}}$ for different MIMO configurations

		MMSE-MRC			MMSE-IRC		
		1 Rx	2 Rx	4 Rx	1 Rx	2 Rx	4 Rx
UMi	1 Tx	57.8%	63.2%	72.4%	57.8%	56.6%	59.1%
	2 Tx	63.7%	68.8%	76.0%	63.7%	63.8%	66.0%
	4 Tx	66.1%	71.0%	76.6%	66.1%	67.4%	68.7%
InH	1 Tx	66.0%	81.2%	81.5%	66.0%	75.9%	78.6%
	2 Tx	69.8%	76.2%	81.7%	69.8%	74.9%	73.4%
	4 Tx	72.0%	76.8%	81.0%	72.0%	75.8%	77.3%

points and, especially in the InH scenario, causes higher losses in our simulations because the restriction to 64-QAM modulation (row e) in Table 2.2) is not so significant at 10.4 dB. In addition, our approach also captures losses from suboptimal link adaptation due to the time-varying nature of the channel and the scheduling in neighboring cells. On the other hand, we do not model the codebook-based precoding inefficiencies as discussed above.

Overall, we observe that in the current LTE standard (or even in general in the current wireless systems engineering state of the art) still a lot of the theoretically possible information rate is wasted. However, increasing the efficiency will be difficult because some factors cannot be changed (e.g., guard bands), others are hard to reduce (e.g., control signaling), and many factors like pilot signal overhead need to be carefully traded-off against the transceiver benefits that they enable. Even bigger gains than the ones suggested by Table 5.17 could come from network-wide coordination, which could potentially remove the interference. But for any such mechanism, an even bigger signaling overhead, especially in the uplink direction, would be required.

5.6. Comparing LTE Transmission Modes 1, 2, 3, 4, and 6

In the previous sections, we discussed different LTE RRM schemes like spatial processing at the transmitter and at the receiver, as well as, for example, scheduling algorithms. Along the way, we also introduced corresponding LTE transmission modes and evaluated their performance. However, not all performance results were completely comparable since we gradually introduced additional features. In this section, we will now present a comparative performance evaluation of LTE transmission modes conducted in an identical setting. In addition, we will put special emphasis on the interference situation under which a link with a certain transmission mode is operated. Previously, when evaluating a certain transmission mode or RRM scheme, we conducted a simulation where all links in the system were operated the same way. For example, in a TM4 simulation, all users in all 57 cells of an outdoor scenario were served using TM4 so that both the evaluated serving link as well as all interfering links were using TM4.

Table 5.18.: Simulation parameters for the simulations presented in Section 5.6

Parameter	Value
Scenario	UMi and InH
Number of drops	15 for InH, 5 for UMi
Number of MS	on average $N_{\text{MS}} = 25$ for InH and $N_{\text{MS}} = 10$ for UMi; evaluating only those 10% of users with the serving TM and rank
Subcarrier modeling	12 subcarriers
BS and MS antennas	2×2 and 4×4 , cross-polarized antennas on both sides (configuration E-E) with $\lambda/2$ spacing in case of four antennas
Receiver type	MMSE-MRC and MMSE-IRC with perfect IPN estimation
Transmission modes	1, 2, 3, 4, and 6 with different maximum number of layers
Link adaptation	with outer-loop link adaptation (OLLA)
Subband sizes	subband CQI (i.e., 3 PRBs per subband) and wideband PMI
Scheduler	proportional fair (PF)

This is in line with the common practice in the literature (see, e.g., the system-level results presented in [36.814]). However, we will see that the performance results for a certain transmission mode can differ materially depending on what transmission modes and ranks are used in neighboring cells.

Table 5.18 lists the simulation parameters used for the comparison in this section. To limit the amount of results, we pick the urban micro (UMi) and indoor hotspot (InH) scenarios for the performance evaluation in a 2×2 MIMO setting. In Section 5.5.2 we have seen that higher transmission ranks are relevant primarily in indoor hotspot scenarios. Thus, we restrict the discussion of 4×4 MIMO results to the InH scenario.

To better control the interference situation that each transmission mode encounters, we form two randomly selected groups of users. We use only 10% of the users for the statistical evaluation and specify a *serving* transmission mode, whose performance is evaluated. Independent of the considered 10% of users, we serve the remaining 90% of the users with an *interfering* transmission mode. In addition to evaluating all possible combinations of serving and interfering transmission modes, we also limit the maximum number of layers that can be used in order to further control the spatial properties of the interference. Note that with the random classification of users, a limited number (i.e., almost 10%) of interfering links also use the *serving* transmission mode and thus create corresponding inter-cell interference. Nevertheless, the majority of interference sources is configured to use the *interfering* transmission mode.

In a 2×2 setup, we consider 25 possible combinations of serving and interfering transmission schemes as shown in Fig. 5.22 for the UMi scenario and in Fig. A.8 in the Appendix for the InH scenario. In a 4×4 setup, we have 81 different combinations because for TM3 and TM4 we can support up to four instead of two layers. The results for the InH scenario with a 4×4 configuration are given in Fig. 5.23 for the MMSE-MRC receiver and in Fig. 5.24 for the MMSE-IRC receiver. We distinguish between the per-cell system spectral efficiency and the cell-edge user throughput performance.

5. Performance of Fundamental RRM Schemes

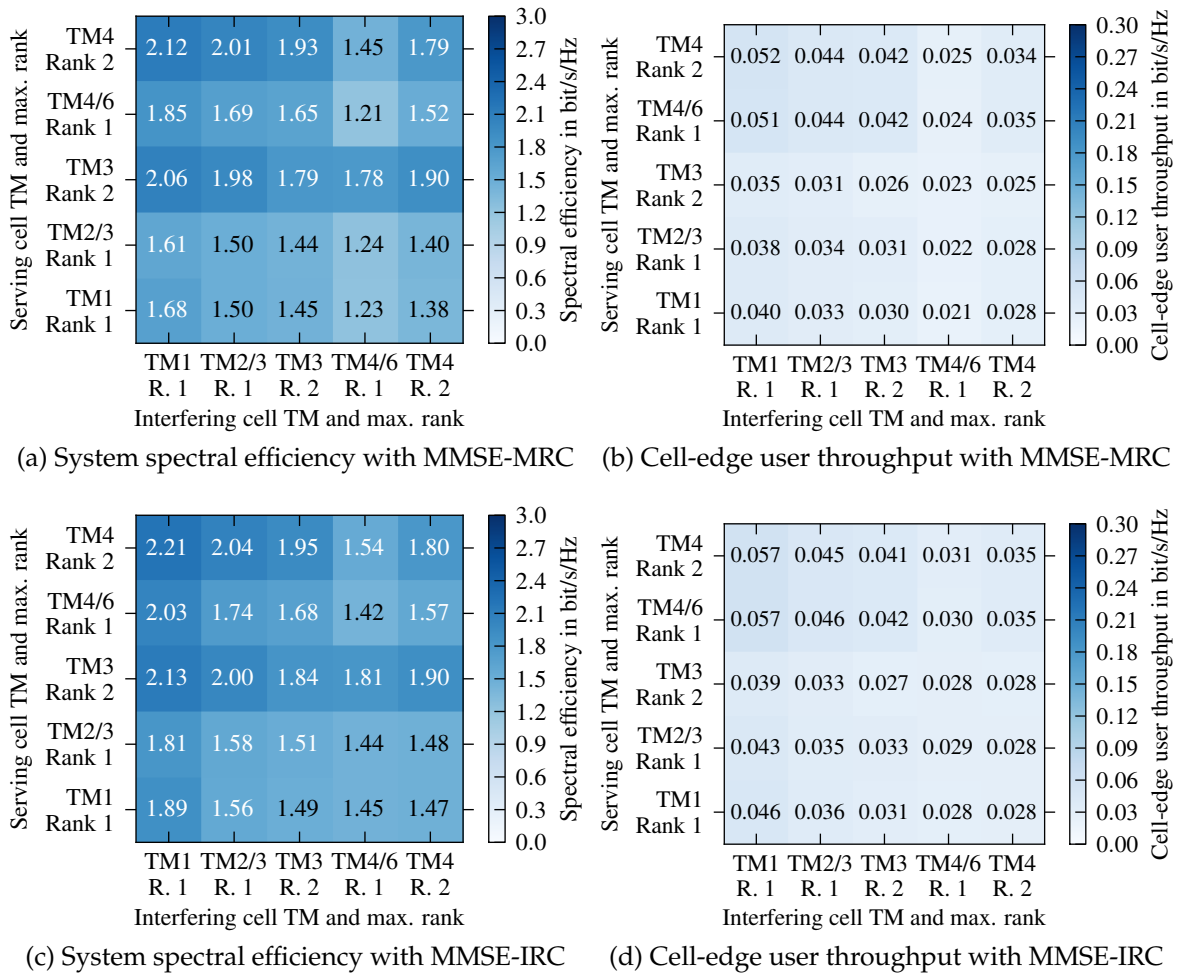


Figure 5.22.: System and cell-edge user spectral efficiency in the UMi scenario with a 2×2 antenna configuration (E-E), evaluating those 10 % of users with TM and max. rank according to the y-axis while the other 90 % of users cause interference with TMs and max. ranks as indicated on the x-axis

The results tables are color-coded according to the colorbar next to the table. Note that the cell-edge results are normalized per-user throughput values so that they are scaled-down by the average number of users per cell (i.e., 10 for UMi and 25 for InH) compared to the system spectral efficiency values. Two transmission mode and rank combinations are consolidated because they are identical: TM4 operation limited to a single layer is identical to TM6 (rank-1 precoding) and TM3 with a single layer is identical to TM2 because for rank-1 operation SFBC transmit diversity is used in TM3. Note that the 2×2 setup with cross-polarized and thus uncorrelated antennas is not well-suited for rank-1 precoding with TM6 or TM4 rank-1.

The diagonal entries of the results tables show the conventional results where the transmission on a serving link is only interfered by links in neighboring cells with the same transmission mode. Apparently, other combinations of serving and interfering trans-

mission mode and rank can yield very different results as indicated by the off-diagonal entries. For example, from the diagonals in the UMi 2×2 setup shown in Fig. 5.22, one could conclude that TM4 with up to two layers performs worse than a simple single-antenna transmission with TM1 in terms of spectral efficiency and cell-edge user throughput, except for the spectral efficiency with MMSE-MRC receivers shown in Fig. 5.22a. However, when comparing apples to apples, that is, when comparing the TM4 performance versus the TM1 performance when both are (mostly) interfered by TM1 transmissions, we observe a clear advantage for spatial multiplexing with TM4 (compare the upper left and lower left corners in all tables in Fig. 5.22). Each table column allows a fair comparison of serving transmission modes against a common interfering transmission scheme.

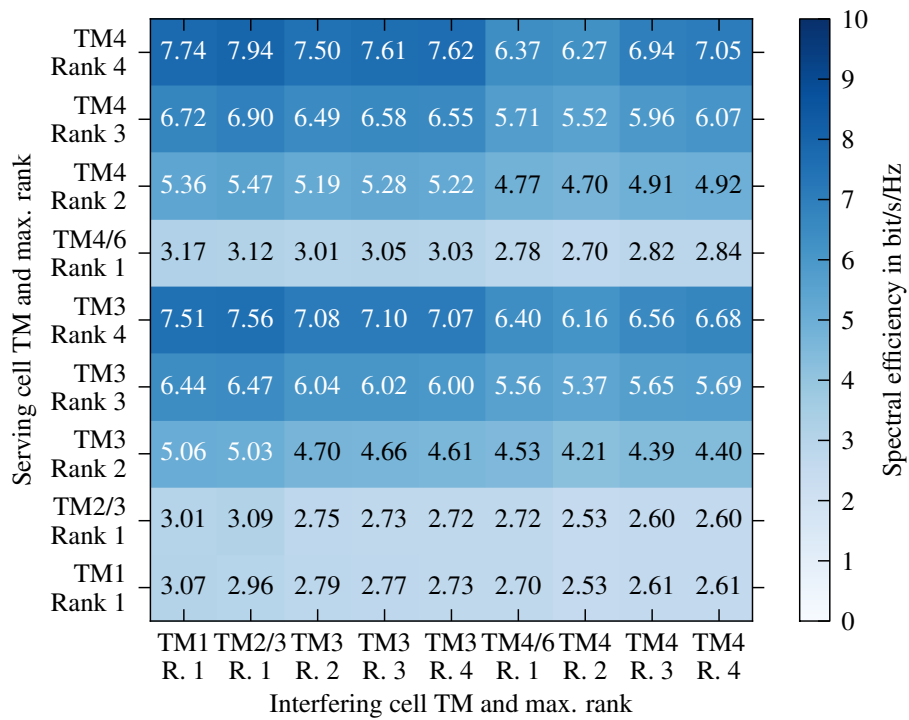
We cannot discuss each and every single result, but we highlight a few general trends: As seen before, transmit diversity with TM2 is, in the low-mobility channel-adaptive scheduling conditions presented here, not advantageous compared to single-antenna transmission with TM1. Regarding the general influence of the spatial interference properties, we can make the following observations: On the one hand, especially the MMSE-IRC receiver prefers single-layer interference (i.e., TM1, or TM6 and TM4 with rank-1) because few dominant interfering layers are easier to suppress. On the other hand, with a higher mix of full-rank (rank-2) interfering transmissions, the flashlight effect is reduced and thus the link adaptation improves, which manifests in a better performance of the MMSE-MRC receiver with full-rank interference compared to rank-1 interference. Also the TM3 interference, which varies per subcarrier, but whose average impact is stable over multiple TTIs, allows better performance of the link adaptation on the serving links compared to TM4 interference.

For the UMi scenario, the cell-edge performance mostly resembles the corresponding system spectral efficiency result. For the InH scenario, however, we observe a significant advantage of the MMSE-IRC receivers when serving and interfering transmissions are limited to a single layer (see Fig. A.8 in the Appendix). In that case, the MMSE-IRC receiver with 2 Rx antennas has just enough degrees of freedom to suppress the single interferer in the InH scenario while receiving the single layer from the serving BS.


We observe a similar situation in the InH scenario with a 4×4 configuration. Figure 5.23 gives an overview of the performance with MMSE-MRC receivers, while Fig. 5.24 shows the MMSE-IRC performance results. As already seen in Section 5.5.2, single-layer schemes are limited by the maximum per-layer spectral efficiency of about 3.77 bit/s/Hz. Both TM3 and TM4 can successfully increase spectral efficiency and cell-edge user throughputs with multi-layer transmissions. Again, especially the MMSE-IRC receiver prefers interfering transmissions with few layers. We also observe that TM3 interference is preferred over TM4 interference with the same number of layers, which we explain by a higher predictability and thus reduced flashlight effect with TM3. According to the higher degrees of freedom with 4 Rx antennas, the sweet spot for high cell-edge throughput with MMSE-IRC interference suppression is now shown by the deeper-blue entries in Fig. 5.24b for combinations where the number of serving and interfering layers adds up to a total maximum of four layers.

Summarizing this comparison of LTE transmission modes 1, 2, 3, 4, and 6, we highlight

5. Performance of Fundamental RRM Schemes




(a) System spectral efficiency




(b) Cell-edge user throughput

Figure 5.23.: System and cell-edge user spectral efficiency in the InH scenario with a 4×4 antenna configuration (E-E) with an MMSE-MRC receiver, evaluating those 10% of users with TM and max. rank according to the y-axis while the other 90% of users cause interference with TMs and max. ranks as indicated on the x-axis

5.6. Comparing LTE Transmission Modes 1, 2, 3, 4, and 6



(a) System spectral efficiency with MMSE-IRC



(b) Cell-edge user throughput with MMSE-IRC

Figure 5.24.: System and cell-edge user spectral efficiency in the InH scenario with a 4×4 antenna configuration (E-E) with an MMSE-IRC receiver, evaluating those 10 % of users with TM and max. rank according to the y-axis while the other 90 % of users cause interference with TMs and max. ranks as indicated on the x-axis




Figure 5.25.: Geometry of simplified cellular model with SIR distribution for $\gamma = 4$

that besides all the deployment aspects (i.e., the scenario, antenna array setups, etc.) and RRM methods (i.e., scheduling and link adaptation algorithms, spatial processing options, etc.), also the interference situation has a major impact on the performance of an LTE link. The performance and ranking between transmission modes can differ significantly depending on whether they are evaluated against a common interference backdrop or in a homogenous setting where each link faces interference from links of its kind. Since the BS can choose the transmission mode for each link individually, an intelligent per-link TM selection algorithm has the potential to improve the overall system performance significantly. Besides deployment-dependent aspects like antenna arrays and radio channel characteristics, which the BS knows or can deduce from measurements and feedback, some relevant aspects for TM selection are a priori unknown to the BS. For example, the BS does not know how many antennas the MS has, what kind of receiver it employs, or what the interference situation looks like. However, with the help of the rank feedback, the MS can signal the appropriate number of layers taking its number of Rx antennas, its receiver implementation, as well as the interference conditions at the time of feedback computation into account.

5.7. Area Spectral Efficiency

In the discussion so far, we always looked at the spectral efficiency per cell without focusing too much on the size of the cell. For a cellular operator, who has to cover a certain geographical area, the *area spectral efficiency* [Hat77; AG99] measured in bit/s/Hz/m^2 is a natural KPI. It relates the spectral efficiency of one cell to the area covered by that cell. As we will show in this section, however, the spectral efficiency per cell in a multi-cell deployment is approximately independent of the cell size. This motivates restricting the discussion in this thesis to the spectral efficiency per cell because the area spectral efficiency then scales more or less linearly with the number of cells. Thus, an operator can simply increase its area spectral efficiency by adding more cells as the demand grows if the costs for the additional equipment and

the base station sites are economically feasible. This technique of building out cellular capacity is known as *cell splitting* at least since the days of the Advanced Mobile Phone Service (AMPS) system in the 1970s [Mac79]. In order to show that the downlink cell spectral efficiency is approximately independent of the cell size, we will first consider a simplified model before presenting system-level simulation results.

As we have seen before, the user throughput scales logarithmically with the SINR of the user as long as the link adaptation works properly. Thus, for our simplified model, it is sufficient to show that varying the ISD in a multi-cell system does not alter the SINR CDF. We will show this in the following. We do not consider isolated cells and are only interested in the efficiency of multi-cell systems with cell sizes well below the theoretical maximum, which is on the order of tens of kilometers for LTE [STB11]. Such systems are always *interference-limited*. That is, the inter-cell interference is always much stronger than the thermal noise and the noise figure of the receiver. Hence, we can approximate the SINR with the *signal to interference ratio* (SIR). Figure 5.25 shows a simplified non-sectorized deployment model with base stations located in the center of the cells. The cell under consideration in the middle is surrounded by one ring of six interfering cells. Each cell has the radius $R = 1/2$ ISD. We express the location of a user in the cell in polar coordinates (r, ϕ) , where we choose r to be a normalized fraction $0 \leq r \leq 1$ of the cell radius R . We denote the distance of the MS to its serving base station by $d_S(r, R) = r \cdot R$ and the distance to an interfering base station as d_I . With the notation shown in Fig. 5.25, we can derive the distance $d_{I,i}$ to the interferer i by the law of cosines as

$$\begin{aligned} d_{I,i}(r, \phi, R) &= \sqrt{d_S^2 + (2R)^2 - 2d_S(2R) \cos(\phi - i \cdot 60^\circ)} \\ &= \sqrt{r^2 R^2 + 4R^2 - 2rR2R \cos(\phi - i \cdot 60^\circ)} \\ &= R \sqrt{r^2 + 4 - 4r \cos(\phi - i \cdot 60^\circ)}, \end{aligned} \quad (5.10)$$

where i is the index of the interferers that are uniformly distributed around the serving cell. If we further model only the distance-dependent pathloss as the dominant large-scale propagation effect, we can write the received power depending on the distance d as $P_{\text{rx}} \sim P_{\text{tx}} \cdot d^{-\gamma}$, where γ is the pathloss exponent as introduced in Section 2.1.1. We use this relation with identical pathloss exponents for the serving and all six interfering links and can thus write the SIR as

$$\text{SIR}(r, \phi, R) = \frac{P_{\text{rx}}^{(S)}(r, \phi, R)}{\sum_{i=0}^5 P_{\text{rx}}^{(I_i)}(r, \phi, R)} = \frac{P_{\text{tx}} [d_S(r, R)]^{-\gamma}}{P_{\text{tx}} \sum_{i=0}^5 [d_{I,i}(r, \phi, R)]^{-\gamma}} \quad (5.11)$$

$$\begin{aligned} &= \frac{(r \cdot R)^{-\gamma}}{\sum_{i=0}^5 \left[R \sqrt{r^2 + 4 - 4r \cos(\phi - i \cdot 60^\circ)} \right]^{-\gamma}} \\ &= \left[\sum_{i=0}^5 \left(1 + \frac{4}{r^2} - \frac{4 \cos(\phi - i \cdot 60^\circ)}{r} \right)^{-\frac{\gamma}{2}} \right]^{-1} = \text{SIR}(r, \phi). \end{aligned} \quad (5.12)$$




Figure 5.26.: Cell spectral efficiency and area spectral efficiency over ISD in the down-link of a 1×2 UMa deployment with MRC receivers

We observe that the SIR only depends on the position of the user in the cell as expressed by the normalized polar coordinates (r, ϕ) with $0 \leq r \leq 1$. That is, the SIR distribution inside the cell is independent of the cell radius and thus of the inter-site distance. The SIR distribution for pathloss exponents $\gamma = 4$ plotted in Fig. 5.25 suggests that the SIR is even almost independent of the polar angle if the interferers are distributed with an equiangular separation around the considered cell. If the users are uniformly distributed over the area of the cell, we can thus expect approximately identical SINR CDFs for cells with different radii and hence cell spectral efficiencies that are independent of the cell size.

Note that we made a couple of simplifying assumptions in this model. A more realistic model would need to account for sectorized deployments, antenna patterns, different pathloss exponents (e.g., for LOS versus NLOS propagation), shadowing effects, multipath propagation, and consider the effect of RRM schemes. To account for all these standard modeling assumptions, we performed a system-level simulation campaign where we modify the ISD of a UMa deployment scenario. Figure 5.26 shows that in a wide range around the standard 500 m ISD, the cell spectral efficiency stays almost constant. The area spectral efficiency grows accordingly with the smaller cell sizes. Only for very short inter-site distances we observe a decrease of the spectral efficiency because we do not adapt the downtilt of the vertical antenna pattern. The downtilt is optimized for an ISD of 500 m so that it results in poor coverage in the own cell and excessive interference for neighboring cells when applied to smaller cells.

These observations justify our focus on cell spectral efficiency throughout this thesis. But we keep in mind that the areal spectral efficiency can be increased by adding more sites, which results in smaller cells. In a similar way, adding more cell sectors to existing sites also leads to more and smaller cells and thus to a corresponding increase in the area spectral efficiency.

5.8. Summary and Contributions

In this central chapter of the thesis, we provided a holistic and quantitative evaluation of the system spectral efficiency performance of the LTE downlink under various radio resource management schemes and deployment scenario configurations. Most of the individual results are qualitatively known from the literature. However, their analysis is spread over a vast amount of publications, each with different modeling assumptions. The main contribution of this chapter thus lies in the consistent evaluation yielding comparable results in a fully-reproducible way. This is especially important considering that most RRM schemes cannot be judged in isolation due to their interplay and the dependence on scenario and deployment parameters, which we highlighted in this chapter. Taking the BS and MS antenna link configurations from SISO to MIMO as landmarks, we gradually introduced additional RRM schemes and assessed their incremental performance contribution. In particular, we found the following:

Channel-adaptive scheduling, for example a proportional fair scheduler, can significantly increase the downlink spectral efficiency if enough multi-user diversity and accurate-enough channel state information is available. The diversity mainly depends on the number of active users per cell and the number of subbands for which CQI feedback is provided. For example, in the presented UMi scenario, 10 active users and a relatively small number of CQI subbands are sufficient to increase the spectral efficiency by more than 50 % compared to the single user case. The other crucial requirement is that the CQI feedback is provided with a feedback delay within the channel coherence time. The latter is thus an example for a scenario dependency regarding the carrier frequency and user velocity. One way to cope with inaccurate CQI feedback is using a so-called outer-loop link adaptation scheme, which is not a standardized feature. We provided an OLLA algorithm and evaluated its performance. Another aspect of the proportional fair scheduler we demonstrated is that it allows trading-off the average cell spectral efficiency against the cell-edge user throughput.

As the first spatial processing RRM schemes, we evaluated the gain from multiple receive antennas at the MS side. Using either an MRC or IRC linear receive filter, significant gains can be achieved that increase with the number of available Rx antennas. The IRC receiver proves to be vastly superior in interference-limited multi-cell scenarios as it realizes a substantial lead over MRC by suppressing inter-cell interference. Since the adaptive Rx filters rely on serving channel and (for IRC) IPN covariance matrix estimates, we evaluated the impact of estimation errors using the error models introduced in Chapter 4. We concluded that the expected Rx filter degradation due to serving channel estimation errors is small enough to allow sticking to perfect channel estimation for comparison reasons. In the case of IRC performance under IPN covariance estimation errors, already the worst-case assumptions in the Wishart-based error model offer a better performance than IPN-unaware receivers and for most use cases, we can stick to perfect IPN estimation as well. Overall, having more MS Rx antennas for spatial processing is attractive from a network operator perspective because no further changes in the system are needed. However, placing more than two antennas on a handheld device is difficult and adds additional transceiver and baseband costs.

With regard to spatial processing at the BS, we first considered the Alamouti-based SFBC transmit diversity scheme in LTE. Under the standard IMT-Advanced evaluation conditions, it did not appear suitable for increasing the spectral efficiency because its channel averaging effect is detrimental to channel-adaptive scheduling. In contrast, the codebook-based beamforming, known as rank-1 precoding in LTE, appears more interesting. Significant gains can be achieved with two or four Tx antennas at the BS. However, the gains from additional BS Tx antennas are not as big as those from additional MS Rx antennas due to the coarse CSI quantization with codebook-based feedback. Another detrimental system-level consequence of Tx-side spatial processing is the so-called flashlight effect that deteriorates the link-adaptation performance due to a larger interference uncertainty. In preparation for an advanced scheme presented in Chapter 6, we also evaluated the distribution of the chosen PMI codebook entries.

For the evaluation of spatial multiplexing MIMO RRM schemes, we discussed the LTE open-loop (TM3) and closed-loop (TM4) MIMO transmission modes. A comparison of both under different scenario conditions showed that due to lower SINRs, TM4 is advantageous in the outdoor scenarios, while in high-SINR indoor hotspot scenarios no significant difference is observed. Again, the IRC variant of the MMSE receiver substantially outperforms the interference-unaware receiver. But, we also noted that in a TM3 / TM4 deployment, the MMSE-IRC receiver suffers from the higher number of interfering layers so that in some cases, a pure TM1 deployment with MMSE-IRC receivers would achieve higher data rates, especially for cell-edge users. In addition, we demonstrated how the chosen ranks and thus the spectral efficiency depend on the antenna configurations. Cross-polarized antennas allow higher ranks compared to co-polarized ones, but offer less array gain for cell-edge users. Operators thus have to make a scenario-dependent choice when deploying antenna arrays at the BS.

Taking the TM4 spatial multiplexing setup as an example, we also quantified the link-level PHY losses observed in system-level simulations. When compared against mutual information rates, the PHY overhead, linear receivers, and limited-feedback link-adaptation and precoding schemes result in a loss between 57 % and 82 % suggesting that future cellular standards could benefit from further link-level optimizations.

A major contribution of this chapter is the consistent comparison of the spectral efficiency and cell-edge throughput for LTE transmission modes 1, 2, 3, 4, and 6 in different IMT-Advanced scenarios. In particular, we highlighted the dependency on the transmission modes applied in interfering cells. It is common practice in the literature to evaluate the performance of a transmission scheme by applying it to all links in the system. As already noted above for TM3 and TM4 versus TM1, this could lead to the premature conclusion that TM3 and TM4 perform worse than TM1 under some circumstances. When compared to the same baseline (e.g., TM1 interference), we could show that the more advanced schemes perform indeed better.

We concluded the chapter with a justification why we focus on per-cell spectral efficiency rather than area spectral efficiency. Using a simplified model, we showed analytically that in cellular systems the cell spectral efficiency is roughly independent of the cell size, which we confirmed by a full system-level simulation campaign. The cell density is thus primarily an economic factor.

6. Selected Single- and Multi-Cell RRM Methods

In the previous chapter, we introduced and analyzed the main LTE **intra-cell** RRM schemes and their interdependencies. Now, we want to present a few selected advanced schemes. There are two foci of our discussion: First, we consider various *interference management* techniques, which perform resource allocation decisions in an **inter-cell** context considering the mutual interference coupling between cells. Second, we evaluate multi-user MIMO schemes, which better exploit multiple antennas at the base station because spatial multiplexing gains are no longer limited by low expected ranks of individual users.

Figure 6.1 gives a classification of the interference management and MU-MIMO RRM methods discussed here. They differ by the degree to which the coordination is localized to a single cell or performed system-wide (y-axis). Further, the decision update rate (x-axis) varies from almost static with frequency planning to almost instantaneous with highly-dynamic coordinated scheduling. Low-complexity methods can thus be found in the lower or left part of the diagram, whereas high-complexity methods are located in the upper right part. As another metric, the required MS-BS and BS-BS signaling overhead is expressed by lighter and darker colors representing low and high signaling loads, respectively. In this chapter, we will present a selection of low-complexity methods as indicated in Fig. 6.1. The higher-complexity schemes are briefly discussed in the outlook in Section 6.3.

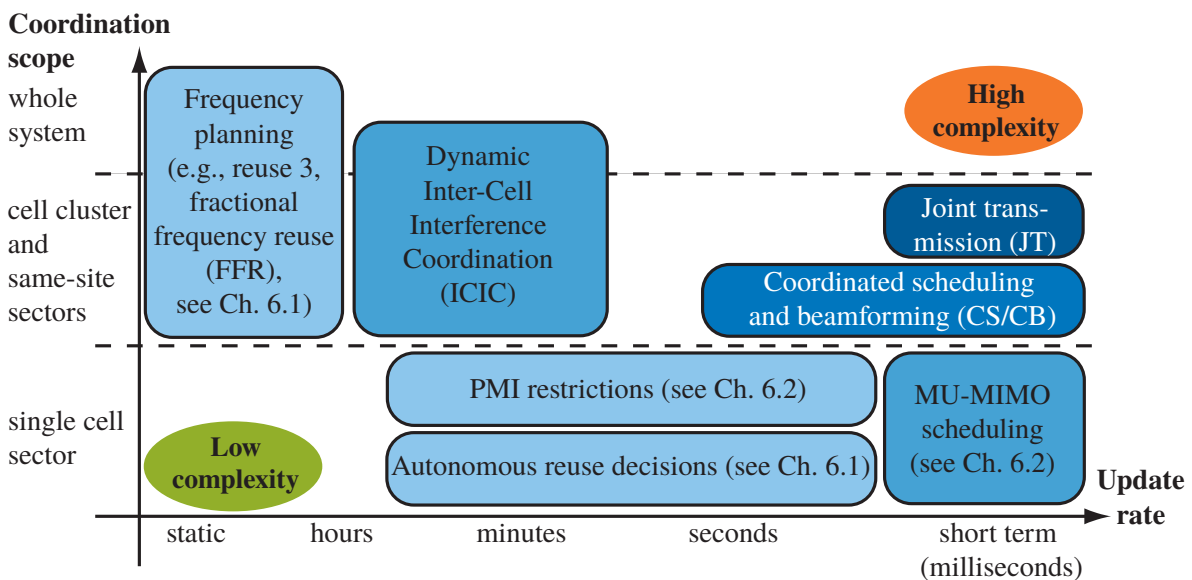


Figure 6.1.: Classification of single- and multi-cell RRM methods w.r.t. coordination scope, update rate, and signaling load (low: light blue, high: dark blue)

6.1. Frequency Reuse Based Interference Management

As we have seen in our simulations presented in Chapters 3 to 5, many so-called cell-edge users achieve only low SINRs and throughputs. The main reason for that is the inter-cell interference in the standard LTE deployments where all cells operate on the same frequency band. In order to improve especially the situation at the cell-edge, but possibly the spectral efficiency overall, various so-called *interference management* techniques have been considered both in the academic literature and in the standardization bodies. Inter-cell interference management was from early on in the focus of the LTE standardization process. In a 2005 technical report [25.814], the 3GPP sketched three potential inter-cell interference management categories in addition to DL beamforming: *randomization* (e.g., by use of cell-specific scrambling like in UMTS), *cancellation* (either by spatial suppression with, e.g., an IRC receiver or by detection and subtraction of interfering signals), and *coordination / avoidance*. Here, we will focus only on the latter. In Chapter 5 we argued that randomization contradicts the intention of channel-adaptive scheduling. The interference suppression aspect of cancellation was already discussed in Chapter 5 and the true cancellation hinges on knowledge about the interfering signal, as we will briefly discuss in the outlook at the end of the chapter.

Inter-cell *interference avoidance / coordination* entails that neighboring cells restrict their resource usage in the time and / or frequency domain in a more or less coordinated way. A cell could either avoid using a resource completely or only use it with reduced power [25.814] in order to lower the inter-cell interference and thus increase the SINR in neighboring cells. The coordination can be updated on different time scales ranging from static planning to updates in every transmission time interval, see Fig. 6.1.

6.1.1. Introduction to Frequency Reuse Schemes

Typically, resource usage restrictions are applied in the frequency domain, which we will assume in the following. But we note that most concepts could be alternatively or additionally applied to the time domain as well.

As we will further discuss in the following section, there is a fundamental trade-off associated with any reuse scheme: stricter resource reuse restrictions per cell lead to less inter-cell interference, better SINRs, and thus higher rates per available spectrum. At the same time, stricter limits reduce the availability of spectrum and thus decrease the overall achievable data rates. As a consequence, any reuse scheme needs to carefully balance these opposing factors.

Overview of Conventional Static Reuse Schemes

Figure 6.2 shows how four conventional static frequency reuse schemes apply transmit power in different parts of the system bandwidth. We assume that the coordination is done per base station site among the three cell sectors of the site.

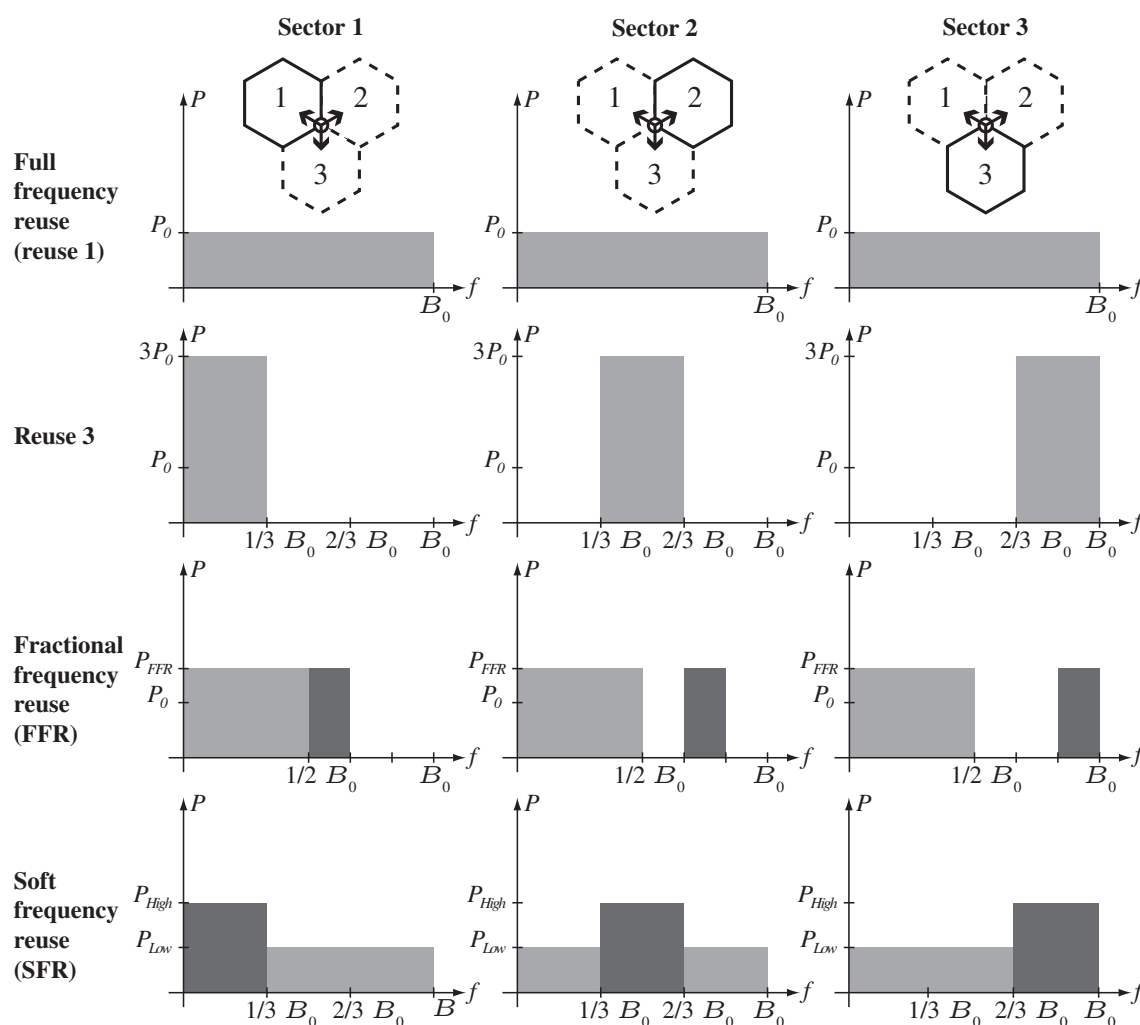


Figure 6.2.: Overview of static frequency reuse schemes

The full frequency reuse or *reuse 1* on the top row of Fig. 6.2 is the conventional scheme that has been assumed throughout this thesis. As we will see, this was for a good reason because it turns out to be advantageous in most relevant use cases. However, it leaves many users with low SINRs and low throughput rates.

The next row shows a *reuse 3* system meaning that the same frequency subband is only reused in every third cell or, expressed alternatively, only one third of the bandwidth is available in each cell. A frequency reuse of three or higher was common in legacy cellular systems. For example, in the circuit-switched GSM system a minimum SINR target was needed to support voice communications. The resulting typical frequency reuse in the original GSM system was 9–12 [FMMO99]. With link-adaptation in the GSM extensions GPRS and EDGE and when employing fractional loading as well as frequency hopping techniques, a reuse 3 scheme was seen as most beneficial [FMMO99; WMJV96]. Overall, a reuse 3 improves the SINRs throughout the cell significantly. However, many users already have good SINRs and cannot benefit enough compared to the drastic reduction in available bandwidth, as we will see.

The *fractional frequency reuse* (FFR) scheme, originally introduced under the name *reuse partitioning* by Halpern in 1983 [Hal83], tries to offer the best of both worlds by balancing a somewhat reduced reuse with aiding only those users who need it. As shown in Fig. 6.2, the spectrum is divided into a part that is commonly used in all cells and into a second part that is coordinated between cells, for example, in a reuse 3 fashion. The idea is that, for users in the “cell center”, no reuse restrictions are needed because the signal from the close-by serving base station is much stronger than the interfering signals. For the so-called “cell-edge” users, however, a reduced frequency reuse is coordinated between neighboring cells in order to lower the inter-cell interference. There are two RRM parameters that can be adjusted. The ratio ω of the total system bandwidth that is available to what we call the *exclusive* part of the frequency spectrum for which a lower reuse is enforced (e.g., $\omega = 1/2$ in Fig. 6.2). And the ratio of users who are assigned to the corresponding part of the spectrum, which we denote with ρ . We will discuss the choice of both in Section 6.1.2 below.

Finally, the *soft frequency reuse* (SFR) scheme, pictured on the bottom row of Fig. 6.2, attempts to overcome the drawback associated with not using some parts of the precious frequency bandwidth. It uses the whole bandwidth in all cells but boosts the transmit power in the coordinated exclusive part of the spectrum compared to the reuse 1 part in order to improve the SINR of the users assigned there. It also means that users in the reuse 1 part of the spectrum experience stronger interference from the overlapping boosted parts in the neighboring cells. The SFR scheme was introduced by Huawei in the LTE standardization process [R1-050507; R1-050841] and enjoyed considerable attention in the literature, see for example [BGW09; XLH07]. However, the final LTE standard makes it unattractive to employ an SFR scheme: The standard requires the pilot power of the CRS resource elements to be uniform across the system bandwidth in order to facilitate better channel estimation. A BS can thus only increase the PDSCH power in the exclusive zone by a user-specific RRC configuration defining the offset between PDSCH and CRS power levels [36.213]. Since RRC configurations are considered to be semi-static, the BS scheduler cannot dynamically switch a user from one zone to the other. This reduces multi-user diversity and complicates the load balancing between the two bandwidth zones. For that reason, we do not consider SFR further in this thesis.

Reuse-Based Inter-Cell Interference Coordination (ICIC) in the LTE Standard

Most static schemes, including the first three schemes introduced above, do not need to be standardized for a BS vendor or network operator to employ them. One drawback of static schemes, however, is that reuse restrictions are applied without considering the cell load (i.e., in terms of active users and their generated traffic) so that a highly loaded cell might limit its capacity by enforcing restrictions even though no neighboring cell could benefit because the rest of the system is unloaded. In order to facilitate dynamic interference management schemes, the 3GPP standardized a few simple signaling mechanisms that report interference situations with a per-PRB granularity by using the X2 interface [36.420] between base stations. Two schemes were standard-




Figure 6.3.: Required SINR improvement to achieve equal mutual information rate with reuse 3 compared to reuse 1 depending on long-term SINR

ized [36.423] for the uplink: an interfered BS can complain about too-high interference on certain PRBs by means of an *overload indicator* (OI). Conversely, a BS can warn its neighbors about PRBs with significant uplink interference by sending a *high interference indicator* (HII). For the downlink, a BS can announce its intention to limit the DL transmit power on some PRBs by signaling a *relative narrow band transmit power* (RNTP) indicator [36.213; 36.423]. Note that the standard only defines the signaling but not what the reaction of the receiving BS should be.

6.1.2. Performance Evaluation of Fractional Frequency Reuse in LTE

In this section, we will evaluate the performance of FFR in the LTE downlink. We limit the discussion to static FFR schemes, which is not a restriction because our standard evaluation assumptions throughout the thesis assume users with full buffer traffic in all cells so that no load imbalances requiring dynamic schemes can occur. Before conducting normal system-level simulations, we first further examine the fundamental interference–reuse trade-off and perform preliminary simulations to study which users might benefit from a lower frequency reuse.

Interference Avoidance versus Resource Reuse Trade-Off

Figure 6.3 sheds some light on the interference–reuse trade-off by showing how much SINR gain a reuse 2 or 3 (at $1/2$ or $1/3$ of the bandwidth, respectively) needs to achieve in order to offer the same data rate as the full frequency reuse. Depending on the reuse 1 (geometry) SINR, the required gains range from about only 3 dB, for the reuse 2 at low geometry values, to more than 30 dB if the reuse 1 SINR is above 15 dB. We derive these curves of equal mutual information rate (see Section 2.1.2) under the assumption of a SISO link by requiring that the reuse 1 rate and reuse R rate are identical, where

the reuse scheme achieves a coordination SINR gain Δ_{SINR} but only has a fraction $1/R$ of resources available:

$$\log_2(1 + \text{SINR}) \stackrel{!}{=} \frac{1}{R} \log_2(1 + (\text{SINR} \cdot \Delta_{\text{SINR}})) \quad (6.1)$$

$$\Leftrightarrow \Delta_{\text{SINR}} = \frac{(1 + \text{SINR})^R - 1}{\text{SINR}} \quad (6.2)$$

Since considerable SINR gains are required to reach or even surpass reuse 1 rates, we now examine if these gains can be achieved in the IMT-Advanced scenarios and which users benefit from a frequency reuse. By identifying the group of users that benefit from a reduced reuse, we can design a FFR scheme where only those users are assigned to exclusive resources. We choose the reuse 3 scheme because it fits well with the sectorized deployments allowing to use different bands in each cell sector, see Fig. 6.2. With a reuse 3 not only the co-located sectors use different bands, but also all neighboring cell sectors in the regular hexagonal tessellation of the plane use disjoint bands so that a good spatial reuse separation is achieved.

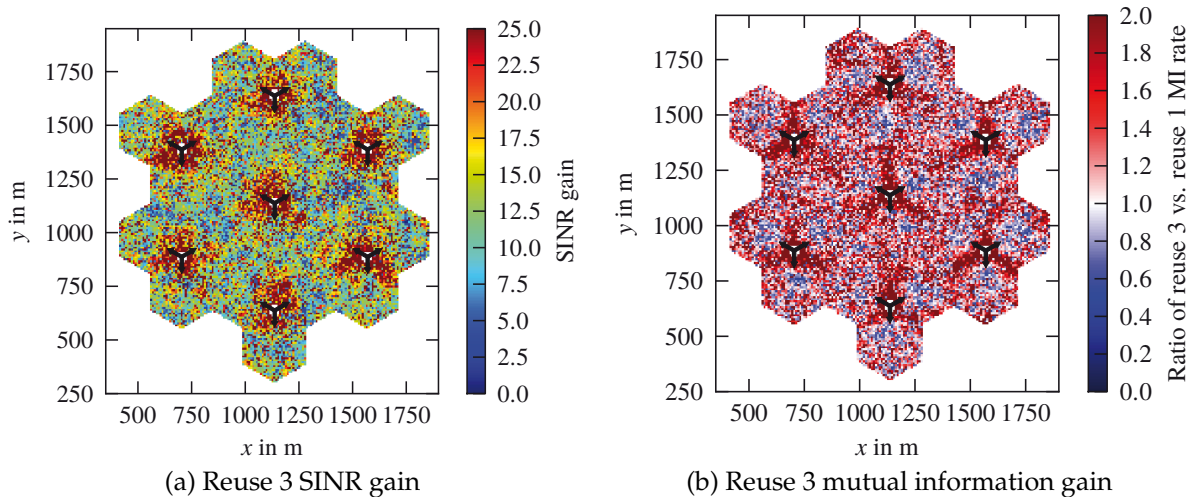


Figure 6.4.: Geographical distribution of the SINR and MI gain in the UMa scenario comparing reuse 3 to reuse 1 with a SISO configuration and RR scheduling

Figure 6.4 shows how the reuse 3 gains (or losses) are distributed over the area of a 21-cell UMa scenario assuming a round-robin scheduling. The scenario is the same as in Fig. 3.9 on page 78, where the geometry over the scenario area is plotted as a reference. According to Fig. 6.4a, the largest SINR gains are not obtained far-away from the BS (i.e., at the “cell-edge”) but rather close to the base station where the otherwise dominantly interfering co-located cell sectors are muted by the reuse 3 coordination. Whether these high gains near the BS or the lower gains farther away lead to a rate benefit from reuse 3, depends on the reuse 1 SINRs at those locations (see Fig. 3.9b). The result is visualized in Fig. 6.4b where we plot the ratio of the reuse 3 versus reuse 1 rate. Blue colors indicate a loss, while red colors indicate a gain from reuse 3. The figure suggests that gains from reuse 3 are obtained at the “cell-sector-edge” rather than in




Figure 6.5.: Distribution of the per-user reuse 3 versus reuse 1 MI gain according to equation (2.27) in the UMi scenario depending on antenna configuration and scheduler; dots indicate scenario-wide average of reuse 3 gain

the middle of the cell. The overall impression is that, for the simple SISO configuration and considering the mutual information rate, more users benefit than suffer from a reuse 3. In the following, we will investigate if more advanced RRM schemes (e.g., PF scheduling and multiple Tx or Rx antennas) change the situation.

Dependency of Reuse 3 Gain on Other RRM Schemes

As shown in Chapter 5, the round-robin scheduling and 1×1 antenna configuration assumed above is not a reasonable selection of RRM schemes. In Fig. 6.5, we compare the scenario-wide distribution of per-user mutual information rate gains from a reuse 3 compared to a reuse 1. Note that we compute the MI rates according to (2.27) on page 23 where we modify the interference plus noise covariance matrix \mathbf{R}_{IPN} to fit the reuse 3 or reuse 1 interference setup. The gains are computed on the PRBs allocated by the corresponding schedulers and, for multiple Tx antennas, considering the chosen TM4 precoders. However, no link adaptation implications can be captured. In contrast to Figs. 3.9 and 6.4, here and for the subsequent system-level simulations, we consider a UMi scenario because the higher speeds in the UMa scenario preclude channel-adaptive scheduling gains (see Section 5.1.2).

In line with the visual impression from Fig. 6.4b, we observe that slightly more than half of the users in the SISO setup with RR scheduling benefit from a reuse 3. Also, the scenario-wide average gain (indicated by the dot on the black CDF curve in Fig. 6.5) shows a small benefit of reuse 3 operation. However, when more advanced RRM schemes are considered, the ratio of users benefiting from reuse 3 shrinks and the overall gains are negative. Even with the SISO setup, a reuse 3 is less advantageous when proportional fair scheduling is used because channel-adaptive scheduling yields a better post-scheduling SINR distribution where reuse 3 gains are less beneficial. The

Table 6.1.: Simulation parameters for Figs. 6.5 and 6.6

Parameter	Value
Scenario	urban micro (UMi)
Number of drops	10
Number of MS	on average $N_{MS} = 10$ (for Fig. 6.5) and $N_{MS} = 20$ users (for Fig. 6.6) per cell moving with scenario-default speed
BS and MS antennas	1×1 , 1×2 , 4×2 , and 4×4 antenna configurations (BS \times MS) with purely vertical polarization and $\lambda/2$ spacing (config C)
Transmission mode	TM1 and TM4
Receive filter	MRC (single layer) and MMSE-IRC (spatial multiplexing)
Link adaptation	with outer-loop link adaptation (OLLA)
Scheduler	proportional fair (PF) and round-robin (RR), PF history exponent $\beta = 1$ for FFR and reuse 3, and $\beta \in \{0.5, 0.75, 1, 1.5, 2.5\}$ for reuse 1
FFR parameters	exclusive bandwidth ratio $\omega = 1/2$, ser ratio $\rho \in \{0.2, 0.3, 0.4, 0.5, 0.6\}$
Channel estimation	perfect channel matrix and IPN covariance matrix estimation
Feedback	subband CQI and wideband PMI
Overhead	CRS ports overhead according to number of Tx antennas

benefits of reuse 3 further deteriorate when spatial processing on both sides is added to the equation. For example, with a 4×4 setup, almost no user benefits from a reuse 3.

Cell Spectral Efficiency and Cell-Edge Throughput for FFR in LTE

The gain distribution observed in Fig. 6.5 with the simplistic SISO RR scheduling clearly motivates the use of a FFR scheme: By applying a full reuse to that half of the users that does not benefit from a reuse 3, we avoid reuse 3 losses for those users. At the same time, we can serve the other half of users with a reuse 3 because we expect them to benefit. The fact that in this simplistic scenario FFR gains seem possible, probably explains why FFR schemes are so popular in the literature. For more advanced (but more realistic) schemes, we do not expect gains based on the above analysis. We will confirm this assumption with a full system-level simulation that also takes more realistic receiver performance (instead of mutual information rates) as well as link adaptation limitations into account.

As introduced above, there are two parameters that can be chosen for an FFR scheme. The ratio ρ of users who get assigned to the exclusive part of the spectrum and the ratio of resources ω used for that part. Here, we choose $\omega = 1/2$ and vary the ratio ρ of users we assign to the exclusive spectrum. In addition to these parameter choices, also an algorithm is needed to classify the users into the full reuse or into the exclusive zone. Many publications (see, e.g. [GML08]) suggest to classify users based on the distance to the BS, which, in a simplistic deployment model with omnidirectional antennas and without (shadow) fading, corresponds to the received signal power (RSRP). Taking Fig. 6.4 as a reference, the RSRP-based classification seems unjustified because also cell-sector-edge users close to the BS with high RSRP benefit from a coordinated reuse.

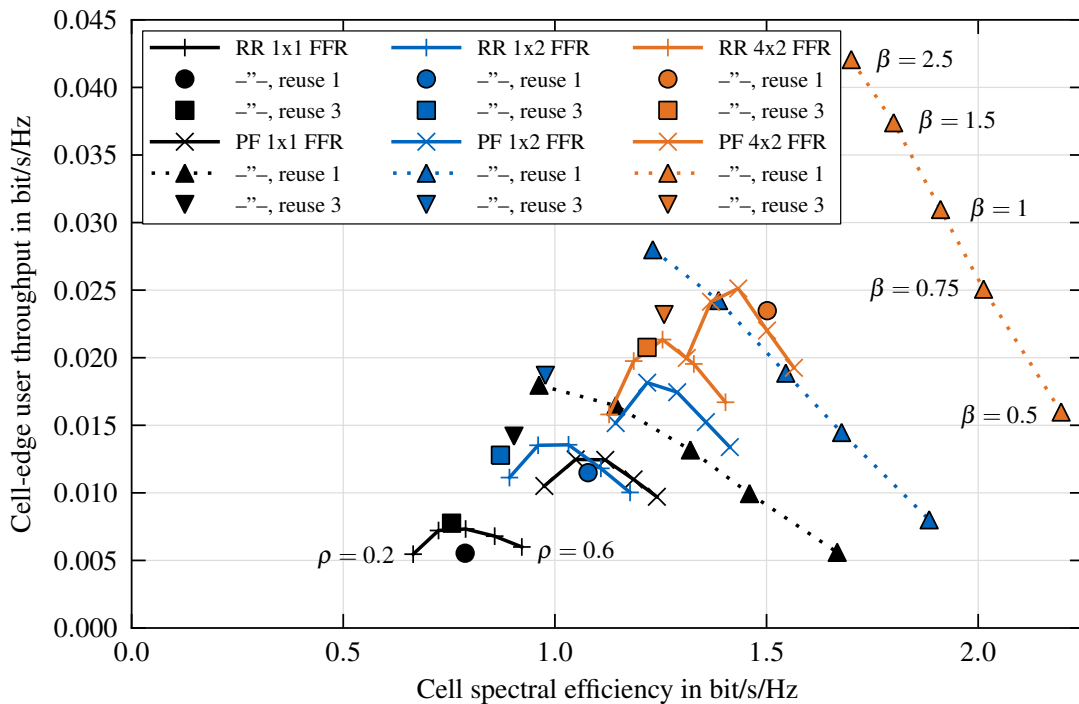


Figure 6.6.: Cell-edge vs. system-spectral efficiency for FFR (user ratio $\rho = 0.2 \dots 0.6$), reuse 1, and reuse 3 with different MIMO schemes and RR and PF schedulers (reuse 1 with PF history exponents $\beta = 0.5 \dots 2.5$); see also Table 6.1

Instead, selecting users with low geometry SINR seems to be a much better-suited criterion because those users have low throughputs and can benefit most from an SINR gain, see Fig. 6.3.

Both variants are implemented in `IMTAphy`, but here we only discuss the latter because it delivers better results as independently confirmed in the literature [Sim07; Iku13]. In our implementation, the $\lfloor \rho \cdot N_{MS} \rfloor$ users having the lowest SINRs are assigned to the exclusive zone. Other FFR variants that classify users according to a fixed threshold SINR (cf., e.g., [Iku13]) risk that parts of the bandwidth are not utilized if no users fall into the corresponding category. In a situation where not all users have full buffers, the classification could be dynamically achieved by adding a bias to the PF metric. This also allows scheduling users in both reuse zones if required by the load situation.

Figure 6.6 shows the cell-edge user throughput versus the cell spectral efficiency obtained by an FFR scheme and by a reuse 3 as well as a reuse 1 configuration with either round-robin or proportional fair scheduling. For the FFR schemes (denoted by solid lines), different user ratios ρ are shown, whereas for the reuse 1 scheme with PF scheduling (denoted by dotted lines) different PF history exponents β are used. With the RR scheduler and a SISO configuration, the FFR scheme, for some values of ρ , achieves modest gains over reuse 1 and reuse 3 in both cell-edge and system spectral efficiency. This accords with our observation from Fig. 6.5 indicating that a reuse 3, applied to the right users, is beneficial. Note, however, that in contrast to the overall reuse 3 MI rate gain over reuse 1 in Fig. 6.5, reuse 3 is inferior in an LTE system-level

simulation because gains in the high-SINR region cannot be realized with the highest-efficiency MCS in LTE. As expected from the discussion of Fig. 6.5, for other RRM scheme combinations the situation changes, and FFR with any choice of parameter ρ is clearly outperformed by a full reuse. When using a PF scheduler for both FFR and reuse 1, already a 1×1 setup yields inferior FFR results compared to reuse 1, and the gap widens when considering Rx- and / or Tx-processing schemes possible with 1×2 and 4×2 antenna setups. One argument often raised in favor of FFR is that the parameterization (i.e., choosing ρ and ω in our case) allows some trade-off between cell-edge and total cell spectral efficiency. However, Fig. 6.6 shows that by tuning the history exponent β in the PF metric (5.1) (see also the discussion in Section 5.1.1), a much better cell-edge versus cell-center trade-off and overall higher throughputs can be accomplished with a reuse 1 scheme. In addition, a reuse 1 is easier to deploy in the field and offers the highest peak data rates.

Frequency Reuse Coordination in the Literature and in Live Networks

The above evaluation paints a devastating picture for static downlink frequency reuse coordination schemes. Still, the literature on FFR and SFR performance is full of contradictory results (see, e.g., [PSuQT13; MLD11] and the references cited therein). One main reason for the diverging results in the literature can be found in different system-level modeling accuracies. As we have shown above, the simplest modeling assumptions (e.g., round-robin, no channel-adaptive scheduling, and SISO links) promise small gains that disappear when operating conditions that are relevant in real-world deployments are considered. A study published by Ericsson [Sim07] was one of the first to point out that, as soon as more than one antenna at the MS is modeled (the standard requires at least two Rx antennas!), reuse 1 outperforms frequency reuse coordination schemes in LTE. Similarly, Ikuno [Iku13] notes that throughput gains with FFR are only achieved for round-robin scheduling but vanish once proportional fair scheduling is modeled. These observations underline the importance of detailed and accurate modeling of the wireless channel and the LTE RRM schemes. Without the efforts that went into *IMTAphy*, we thus might have only been able to derive misleading results regarding the performance of FFR schemes.

As a consequence of their bad performance, downlink frequency reuse coordination schemes are not used—to the best of the author’s knowledge—in any LTE network deployed today. However, in the uplink of some live networks, a coordinated fractional loading of cells in the frequency domain can be observed. As long as the network is underutilized, neighboring cells can avoid creating interference to each other, for example, by starting at different edges of the frequency band when they allocate PRBs to their users. Compared to colliding allocations, this improves uplink SINRs at the BS, allowing the users to transmit with less power. Note that a given cell on its own cannot increase its sum throughput by using less resources because using more resources would always be beneficial. However, a fully-loaded cell can benefit if its underutilized neighbor cells coordinate their frequency resource usage because that leads to less fluctuating interference, which improves the link adaptation.

6.1.3. Dynamic Frequency Reuse Coordination — An Excursion

We have seen that static frequency reuse coordination does not promise spectral efficiency gains in macro cell wireless systems because reusing the full spectrum is always beneficial. However, this does not mean that all kinds of reuse coordination schemes are useless. On the one hand, dynamic schemes that adapt to traffic load variations between cells could make many or all resources available in highly loaded cells while cells with low load employ a restricted and coordinated reuse. On the other hand, there can be situations where it is technically infeasible to use all available frequency bands in a cell. In such situations, a reuse coordination between cells offers interference avoidance gains without the need to sacrifice additional spectrum.

Examples for scenarios with highly varying traffic loads and the technical need to restrict spectrum usage can be found in so-called *femtocells*, which are user-deployed small base stations, also called *home eNodeB* in the LTE context [CAG08]. Since a femtocell usually only serves devices from a closed subscriber group (e.g., only devices belonging to the household where it is deployed), the traffic load can be very bursty with long phases during the day when there is no load at all. In addition, a femtocell BS might be forced to select a subset of spectral resources to employ in its cell because it might be power-limited or only able to utilize a limited amount of LTE-Advanced component carriers in parallel. A *component carrier* refers to one of the frequency bands that can be aggregated for transmission per cell using *carrier aggregation* [SPM+12], which is arguably the commercially most important feature of LTE-Advanced (Rel-10). For example, carrier aggregation allows a base station to serve a user by simultaneously transmitting on multiple separate bands (e.g., at 800 MHz, 1.8 GHz, and 2.6 GHz) if the operator owns fragmented spectrum. Since the baseband and transceiver complexity scales with the number of simultaneously supported bands, a low-cost femto base station might be forced to select a suitable subset of bands and would thus want to coordinate the selection with its neighbors.

Femtocells are not deployed in a pre-planned fashion so that severe interference situations between closely deployed femtos as well as between a femto and its macro cell can arise. For the support of so-called *heterogeneous networks (HetNets)* containing high- and low-power cells, an *enhanced inter-cell interference coordination (eICIC)* mechanism was introduced in LTE Rel-10 [PWSF12]. It provides a mechanism for interference coordination between a high-power macro cell and a low-power femtocell. With eICIC, the macro cell protects the transmissions in the low-power cell by not transmitting PDSCH transmissions during explicitly signaled *almost blank subframes (ABS)*. We do not consider such macro-to-femto interference scenarios in this thesis. Instead, in this section we present an algorithm that could be used to achieve a decentralized femto-to-femto interference coordination.

The algorithm we discuss in this section was inspired by a similar method suggested for ad-hoc networks [NC06]. It was first applied to cellular systems and compared to the optimal allocation by the author of this thesis in [EllWiWi07; Ell08]. Originally presented for interference coordination in macro cell networks, our approach has been used in the literature for various purposes. Holtkamp [Hol13] used it in his PhD thesis

to align the *discontinuous transmission* (DTX) phases of base stations, which in essence is a reuse coordination in the time domain in order to save power. Da Costa et al. [dC-CKM10] applied it to the femtocell use case described above, and in [GKP+12] the same research group also suggested to perform interference coordination on a component carrier level in LTE-Advanced.

We consider this section to be an excursion compared to the remainder of the thesis because we do not provide simulation results on the same level as before for two main reasons. First, even though the 3GPP defined femtocell deployment scenarios in [36.814], there is no suitable IMT-Advanced [M.2135] model for this use case. Second, system-level simulation results would be hard to compare to the results presented throughout the thesis due to the very different setup. Thus, we refrained from supporting femtocells in IMT-Aphy.

Motivation for Autonomous Femto-to-Femto Interference Coordination

Figure 6.7 shows the so-called *dual strip* femtocell deployment defined by the 3GPP in [36.814]. It consists of two large buildings that contain 20 individual apartments each. The assumption is that in many apartments (colored gray) a femtocell base station is deployed to improve the in-building coverage of the residents. For the example in this section, we have selected the fixed random femto deployment as shown in Fig. 6.7. One important source of interference to the femtocells is the macro cell BS that covers the area of the buildings. But due to thin indoor walls with an assumed 5 dB attenuation and the very dense deployment of femtocells, the femto-to-femto interference is very likely the dominant source of interference.

Given this potentially high inter-cell interference situation and considering our arguments presented above why a femtocell might not be able to utilize all channels, a restricted and coordinated reuse appears to be a very reasonable operation mode for this femto deployment. However, femtocells are not deployed by the operator whose spectrum they are using so that an advance planning and optimization of the resource reuse is impossible to achieve. Since also traffic demands can change dramatically during the day, a dynamic scheme is needed. Such a dynamic reuse coordination could be facilitated by signaling between the femtocells and a core network entity. Besides the signaling load, this would also require a central controller that would have to solve a hard optimization problem.

In contrast to that, we propose a coordination mechanism that achieves a stable system-wide coordination by allowing each femtocell to choose and update its resources in a selfish way. The main insight is that from a femtocell's point of view, it is irrelevant who creates the interference, it just matters how much interference is created. Our proposed channel selection heuristic is thus rather simple: Each femtocell selects those channels that have the lowest interference level. The femto BS either measures the interference level itself, or relies on feedback from its mobiles. The result of the envisioned coordination process in each cell is a pool of resources that is not too highly interfered. The femto BS can then use this subset of resources for scheduling its users, for example in a channel-adaptive way as discussed in Chapter 5.




Figure 6.7.: Dual strip model with two single-floor houses [36.814] where in 20 out of 40 apartments a femtocell is active

There are two main questions we want to answer in the following: Can such autonomous decision making lead to a stable outcome, and how good will that outcome be? We will answer the first question by using solution concepts from *non-cooperative game theory* [OR94] in order to motivate why the suggested iterative coordination based on autonomous decisions converges to a stable *Nash equilibrium* allocation. For the second part, we will evaluate the quality of the stable solutions by comparing them to the optimal pre-planned allocation, to a random selection, as well as to a fully-colliding selection with maximal interference in all cells.

Game Theoretic Modeling

In our proposed scheme [Ell08], each femtocell $t \in \mathcal{T}$ autonomously selects $D_t \leq F$ out of F total channels depending on its demand. We assume the demand to be fixed, for example, due to technical reasons. The choice of cell t is denoted by the binary vector $\mathbf{x}_t \in \mathbb{B}^F$ where $x_{t,f} = 1$ indicates that channel f is part of the chosen subset in cell t . We assume the *channels* to be orthogonal but do not further specify whether we mean PRBs, groups of PRBs, whole component carriers, or even separate time slots. The use of a resource in cell t' causes interference in all other cells $t \neq t'$ depending on the attenuation between cells t and t' . We denote the interference impact the utilization of resource f in cell t' has on the same resource f in cell $t \neq t'$ by the pathloss $PL_{t,t'}$ assuming equal Tx powers in the system. Since the cells are very small and clearly separated by walls, we do not model individual user positions and reduce the resolution to one interference level per channel and femtocell instead. The resource allocation $\mathbf{x}_{-t'}$ of all cells other than cell t' is defined as $\mathbf{x}_{-t'} := (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{t'-1}, \mathbf{x}_{t'+1}, \dots, \mathbf{x}_T)$ with $T = |\mathcal{T}|$. Using this notation, we can express the utility $U_{t'}(\mathbf{x}_{t'}, \mathbf{x}_{-t'})$ of cell t' as the (negative) sum of the interference powers added in the linear domain on all selected channels:

$$U_{t'}(\mathbf{x}_{t'}, \mathbf{x}_{-t'}) = - \sum_{f=1}^F \left(x_{t',f} \sum_{t \in \mathcal{T} \setminus \{t'\}} PL_{t,t'} x_{t,f} \right), \quad (6.3)$$

where each cell only selects as many channels as it demands (i.e., $D_t = \sum_{f=1}^F x_{t,f}$). The pathloss $PL_{t,t'}$ is independent of f and symmetric between cells (i.e., $PL_{t,t'} = PL_{t',t}$).

A game in the context of game theory models a decision situation in which multiple decision makers face a situation of strategic interdependence when deciding their action to take. Here, the base stations are the decision makers or *players* of the game. Their *action* is to each choose an allocation vector $\mathbf{x}_t \in \mathbb{B}^F$ which maximizes their *utility* U . The utility function $U(\mathbf{x}_t, \mathbf{x}_{-t})$ maps the player's own action \mathbf{x}_t as well as the actions \mathbf{x}_{-t} of all other players to some real value. In general, a finite *normal form game* Γ is denoted by a tuple $\Gamma = (\mathcal{T}, (\mathcal{X}_t)_{t \in \mathcal{T}}, (U_t)_{t \in \mathcal{T}})$ where \mathcal{T} is the finite set of players and $(\mathcal{X}_t)_{t \in \mathcal{T}}$ and $(U_t)_{t \in \mathcal{T}}$ denote possibly player-specific action sets and utility functions, respectively [OR94].

Convergence of Decentralized Coordination Based on Autonomous Decisions

If the resource allocation in a cellular system is left to autonomously deciding base stations, one central question is whether the distributed allocation process stabilizes and reaches a steady state. Especially, conditions in which neighboring cells constantly flip-flop in lockstep between two resources f_1 and f_2 might occur if they decide at the same time with the knowledge of the interference level resulting from their previous decisions. In order to show that the distributed coordination process reaches a steady state under certain conditions, we will first show the existence of a Nash equilibrium before we examine algorithms for reaching such a steady state.

A Nash equilibrium action profile $(\mathbf{x}_t^*, \mathbf{x}_{-t}^*)$ describes a steady state in which no player can gain a higher utility by deviating from his chosen action \mathbf{x}_t^* as long as the other players stick to their actions \mathbf{x}_{-t}^* . At a Nash equilibrium $U_t(\mathbf{x}_t^*, \mathbf{x}_{-t}^*) \geq U_t(\mathbf{x}_t, \mathbf{x}_{-t}^*)$ holds by definition for all players $t \in \mathcal{T}$ and for all $\mathbf{x}_t \in \mathcal{X}_t$. It is natural to assume that a base station t would only want to change its resource allocation from \mathbf{x}_t to \mathbf{x}'_t if the total interference is reduced that way, that is, if $U(\mathbf{x}'_t, \mathbf{x}_{-t}) > U(\mathbf{x}_t, \mathbf{x}_{-t})$. Therefore, the Nash equilibrium resource allocation is stable once it is reached.

We assume that the above game is played multiple times in order to *learn* a good solution [FL98]. In contrast to many classical games, for example, the famous *prisoners' dilemma*, the players are unable to deduce an equilibrium solution just from the knowledge of their opponents' utility functions and possible action sets. The reason is that because apart from the interference level, all resources are equal so that for combinatorial reasons, there are generally a vast number of equally good solutions only differing in the resource numbering. The players thus have to learn which particular set of resources is heavily used by other players in order to avoid these resources and settle for the equilibrium allocation.

For our argumentation, we introduce the following *potential function* $\Phi(\mathbf{x}_t, \mathbf{x}_{-t})$ representing half the global interference power added over all cells t' , interferers $t \neq t'$, and channels f :

$$\Phi(\mathbf{x}_t, \mathbf{x}_{-t}) = -\frac{1}{2} \sum_{t' \in \mathcal{T}} \sum_{t \in \mathcal{T} \setminus \{t'\}} \sum_{f=1}^F \text{PL}_{t,t'} \mathbf{x}_{t,f} \mathbf{x}_{t',f}. \quad (6.4)$$

Now we can make the crucial observation that, if a single player t' unilaterally adapts its resource usage by switching from a channel f with higher interference to a channel

f' with lower interference, his utility improvement $U_{t'}(x'_{t'}, x_{-t'}) - U_{t'}(x_{t'}, x_{-t'})$ exactly equals the change in the global potential $\Phi(x'_{t'}, x_{-t'}) - \Phi(x_{t'}, x_{-t'})$. By no longer using resource f , cell t' does not suffer anymore from the sum $\sum_{t \in \mathcal{T} \setminus \{t'\}} \text{PL}_{t,t'} x_{t,f}$ of individual interference contributions coming from all other cells $t \neq t'$ that also use channel f . At the same time, cell t also ceases to cause its interference contribution $\text{PL}_{t,t'}$ to all these other cells $t \neq t'$. The analog is true for the channel f' that t' switches to. Overall, the difference in total interference that cell t' causes and suffers from is exactly identical because the pathloss and thus the interference impact is assumed to be symmetric in our model. Consequently, the total global interference changes by twice the utility difference $U_{t'}(x'_{t'}, x_{-t'}) - U_{t'}(x_{t'}, x_{-t'})$ observed in cell t' because the sum of utility improvements observed by all other cells is equally large. Thus, by definition of $\Phi(x_t, x_{-t})$, the global potential improves by the same amount as the individual cell's utility.

The existence of a global potential function makes the game a so-called *potential game* [MS96] for which a Nash equilibrium always exists. To understand this, we assume that all players act sequentially choosing their actions one at a time after being able to observe the outcome (in the form of interference) resulting from the other players' resource allocation actions. Then, if a player t has a *better response* x'_t that yields a higher utility (lower total interference), he can increase its utility by $U_t(x'_t, x_{-t}) - U_t(x_t, x_{-t})$ which exactly equals the corresponding increase in the game's potential. Note that we assumed all other players' allocations x_{-t} to remain constant while player t adapts. Now, if some other player t' performs a better response to the new interference situation, he will further increase the global potential Φ while improving his own utility. The selfish behavior of the players never makes them adapt a resource allocation yielding a lower utility so that the potential function is monotonically increasing in each round. This effectively prevents circles in the distributed allocation process avoiding that a global allocation profile (x_t, x_{-t}) is reached twice. Since the number of players as well as the size of their action sets are finite, the sequential improvement process has to reach a state in which no player is able to improve his utility by unilaterally changing his action from x_t to x'_t assuming all other actions as constant. This is the definition of a Nash equilibrium.

In order to benefit from the above convergence result, the base stations as the players of the resource allocation game only have to implement a very simple resource allocation algorithm that consists of switching to the least interfered channels. Only local information regarding the interference levels on all resources is needed and the computation complexity is linear in F as only the D_t currently best resources have to be identified. The algorithm behaves deterministically but requires that only one player at a time adjusts to the current interference conditions once they have been observed by measurements.

Provable convergence is a very desirable property, but the required strictly sequential adaptation is problematic for two reasons. First, a suitable serialization mechanism would be required to prevent two femtocells from adapting simultaneously. Second, the adaptation process would become too slow for larger cellular networks. To overcome the scalability problems of strictly sequential adaptations, we implemented

a second algorithm that allows for simultaneous decision making by introducing a p -persistence feature [NC06]. If a femtocell identifies an unused channel with less interference than a currently used one, it switches to that resource only with probability p . The rationale behind this is to prevent neighboring cells, which have similar interference conditions, to always switch simultaneously to the same resource thus causing high mutual interference. Of course, the probabilistic element prohibits proving general convergence because the Bernoulli trial might fail indefinitely long. However, as we will see below, the p -persistent algorithm performs well and exhibits better convergence properties than the sequential one.

Interference Coordination Performance Compared to Benchmarks

Using the femtocell deployment shown Fig. 6.7, we conducted simulations to evaluate the convergence behavior of the sequential and of the simultaneous p -persistent adaptation algorithm. Table 6.2 summarizes the deployment parameters and also defines the femtocell pathloss model that was used. As noted above, we do not perform our usual kind of system-level simulation and only examine how the two schemes can reduce the system-wide average interference per selected channel. We assume an identical reuse of $D_t = 10$ out of 20 channels in all cells yielding an effective reuse 2. Note that the algorithm can also cope with arbitrary demands per cell.

Table 6.2.: Overview of femtocell deployment and pathloss model parameters

Parameter	Value
Deployment	3GPP dual strip model with two single floor houses consisting of 20 apartments each, see Fig. 6.7
Number of cells	20 femtocells randomly placed in 40 apartments, see Fig. 6.7
Channels	20 in total, out of which 10 are used per femtocell (i.e., effectively a reuse 2)
Tx power	0 dBm per channel
Pathloss model	$PL = \max [38.46 + 20 \cdot \log_{10}(d), 15.3 + 37.6 \cdot \log_{10}(d)] + PL_{\text{WallIn}} + PL_{\text{WallOut}}$ in dB, where d is the distance between cell centers; adapted from [36.814, A.2.1.1.2-8] without additional pathloss inside apartment
PL_{WallIn}	$N_{\text{WallIn}} \cdot 5$ dB, where N_{WallIn} is the number of inside apartment walls separating two femtocells
PL_{WallOut}	40 dB if two femtocells are in different buildings, 0 dB otherwise

In order to judge how large the coordination gains achieved at the Nash equilibria are, we use several benchmarks. The *worst case* is represented by an allocation where all cells use the same channels (e.g., the first 10 out of 20) and create maximal inter-cell interference. With a *random selection* of channels, a significant interference avoidance and averaging gain can be realized, as we will see. We also use a random selection in each cell as the starting point for our coordination game and show that further interference improvements are possible by autonomously coordinating cells. Finally, the *optimal* system-wide channel allocation serves as the lower interference bound.

We obtain the optimal allocation in the example deployment realization shown in Fig. 6.7 by employing a classical optimization approach. By stating the interference coordination problem as an *integer linear program* (ILP), we can use standard optimization software like *Gurobi* [Gur13] to find the optimal solution. The ILP is given in equations (6.5) to (6.7). The objective again is to minimize the global total interference (6.5) under the constraint (6.6) that the external demands D_t are fulfilled in every cell. The pair constraint in (6.7) allows writing the problem as an integer *linear* program by avoiding the multiplication of the decision variables $x_{t,f}$ and $x_{t',f}$.

$$\text{minimize} \quad \sum_{t' \in \mathcal{T}} \sum_{t \in \mathcal{T} \setminus \{t'\}} \sum_{f=1}^F \text{PL}_{t,t'} m_{t,t',f} \quad (6.5)$$

$$\text{subject to} \quad \sum_{f=1}^F x_{t,f} = D_t \quad \forall t \in \mathcal{T} \quad (6.6)$$

$$m_{t,t',f} \geq x_{t,f} + x_{t',f} - 1 \quad (6.7)$$

$$\text{optimization variables} \quad x_{t,f}, m_{t,t',f} \in \mathbb{B}$$

$$\text{constants} \quad \text{PL}_{t,t'} \in \mathbb{R} \text{ and } D_t \in \mathbb{N}$$

Figure 6.8 shows the results of our simulations. In Fig. 6.8a we observe that in 25 exemplary simulation runs, both algorithm variants converge to Nash equilibria (indicated by \times marks) that reduce the average interference by several decibels compared to a random start allocation. The simultaneously updating p -persistent scheme converges faster due to the simultaneous decisions and is also seen to reach better Nash equilibria. We have experimentally found $p = 1/7$ to be a good trade-off for p between non-convergence ($p = 1$) and slow convergence (small p). The advantage of the sequential decision scheme is that (in line with our potential-based argument) the system-wide average interference is monotonously decreasing, while for the simultaneous scheme a few short increases are observed.

The results from 10 000 simulation runs shown in Fig. 6.8b confirm the impression from Fig. 6.8a that the simultaneous scheme achieves lower equilibrium interference levels. From the detailed results in Table 6.3, we observe that after 250 iterations, both variants have always converged to an equilibrium. Compared to the worst case channel allocation, already the interference-averaging random allocation achieves an interference reduction of almost 3 dB. On top of that, the autonomous coordination allows an additional gain of 4.5 to 5.5 dB on average. Depending on the SINR operating point, the achieved coordination gain over the worst case (reuse 1) allocation achieved with the autonomously organized reuse 2 might be big enough to offer an overall spectral efficiency improvement, compare Fig. 6.3. If for one of the various reasons discussed above (e.g., a limit on the component carriers), a cell cannot use all channels anyway, a coordinated reuse is always beneficial.

We also note that the distributed schemes are not able to come close to the optimal interference level that a centralized algorithm could enforce. Similar to a concept used in *algorithmic game theory* [NRTV07], we call this experimentally-derived gap the *price*

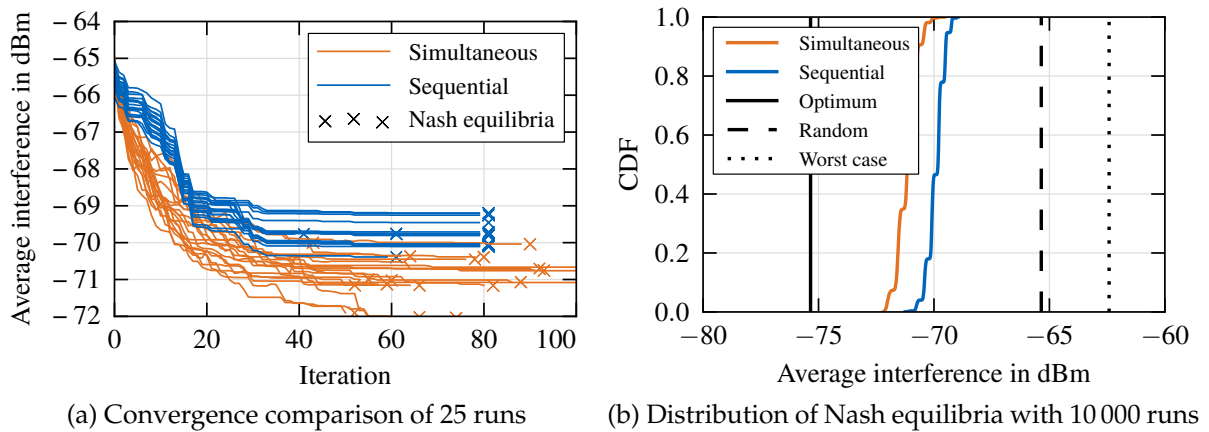


Figure 6.8.: Convergence to Nash equilibria and comparison of equilibrium efficiencies to benchmarks for the sequential and p -persistent schemes

Table 6.3.: Efficiency of Nash equilibria compared to benchmarks after 250 iterations

	Sequential decisions	Simultaneous p -persistent decisions
Total games played	10 000	10 000
Nash equilibria reached	10 000	10 000
Iterations needed (median)	81	75
Worst case allocation		-62.41 dBm
Random allocation		-65.35 dBm
Worst Nash equilibrium	-68.93 dBm	-69.51 dBm
Median Nash equilibrium	-69.80 dBm	-71.13 dBm
Best Nash equilibrium	-71.24 dBm	-72.18 dBm
Optimal allocation		-75.34 dBm
Price of anarchy	4.1 to 6.4 dB	3.2 to 5.8 dB

of anarchy because it is the opportunity cost that a network operator has to pay if it decides to let selfish femtocells compete in the suggested way instead of forcing them to a centralized solution. One can argue that the gap is significant, but a discrepancy between the optimization and Nash equilibrium solutions is not surprising because they achieve a different notion of fairness. The first approach might limit the utility of some cells in order to allow for a higher total utility, while the latter reaches an equilibrium in which no cell could unilaterally improve its utility.

When presenting the scheme, we made a few simplifying assumptions. For example, once different geographic location of users are modeled, the assumption of symmetric interference impacts between cells is no longer valid. However, in system-level simulations good convergence behavior was still observed, and by using explicit inter-cell signaling, a symmetric setup can be enforced [Ell09a]. Another improvement would be to define the utility in terms of cell spectral efficiency and adapt the effective reuse D_i/F accordingly as discussed in [dCCKM10; Da 12].

6.2. Neighbor-Considerate and Multi-User MIMO Transmission Schemes

In the first part of this chapter in Section 6.1, we have seen that a restriction of spectrum use in the time or frequency domain does, in general, not improve the spectral efficiency of a modern cellular system, except for special deployments. In this section, we thus focus on advanced multi-cell RRM methods in the spatial domain that allow a full reuse in the time and frequency domain. We address two issues that we already observed in Chapter 5: First, utilizing multiple antennas at the BS for precoding / beamforming can improve the SINR on the serving link, but the resulting unpredictable spatial interference structure can have detrimental effects on neighboring cells (see, e.g., the flashlight effect observed with TM6 in Section 5.4.3). To this end, in Section 6.2.1 we will demonstrate how the link-adaptation and thus the system performance can be significantly improved by a simple restriction on the use of beamformers. Second, we address *multi-user MIMO* (MU-MIMO) techniques that help to overcome the *single-user MIMO* (SU-MIMO) rank limitations in outdoor deployments (see, e.g., Figs. 5.20 and A.6) by transmitting multiple layers but to different users. Extending the observations from Section 6.2.1, we will consider a MU-MIMO variant called per-user unitary rate control (PU²RC) in Section 6.2.2, which combines predictable intra- and inter-cell interference with MU-MIMO, and compare it to a zero-forcing MU-MIMO scheme.

In the following two subsections, we will first introduce the above mentioned approaches, discuss the relevant literature, and briefly describe our implementation. The performance evaluation will be presented jointly in Section 6.2.3.

6.2.1. Predictable Beamforming

Remember from our discussion of TM6 in Section 5.4.3 that the unpredictable scheduling of beamformers in interfering cells, which can change on a PRB-to-PRB basis and from TTI to TTI, leads to sub-optimal link adaptation performance and thus lower system spectral efficiency. The main reason for the unpredictability is that the base station scheduler makes a new decision each 1 ms which set of users it schedules and on which PRBs it schedules the selected users. In addition, also the desired PMI per user can change.

One radical solution to this problem is to introduce some kind of coordination or even a centralized scheduler that performs the scheduling jointly for multiple cells removing the uncertainty of the scheduling decisions in the interfering cell. However, such *coordinated multi-point* (CoMP) transmission approaches are very demanding in terms of feedback signaling and benefits shown in system-level simulations depend heavily on modeling assumptions. We will thus only refer to them in the outlook in Section 6.3.3. Instead, we choose a very simple, yet effective, way to improve the interference predictability and with it the system spectral efficiency.




Figure 6.9.: Exemplary beamformer (PMI) to subband mapping with 4 beams per cell

We suggest to stabilize the inter-cell interference situation by refraining from user-specific transmit signal processing. In contrast to Ivrlač et al. [IN07], who in their analysis of the fluctuating inter-cell interference in the MISO downlink suggest to use *transmit diversity* for interference stabilization purposes, we employ different beamformers on different subbands, but we enforce using always the same beamformer on a given subband. Which users are scheduled on which subbands can differ dynamically as before. To neighboring cells, however, the user scheduling is irrelevant because the spatial interference structure per subband is stable and thus predictable. As depicted in Fig. 6.9, different subbands (denoted by different colors) employ a fixed beam for transmission on that subband, which can differ between cells.

Despite the fixed PMI-to-subband allocation, the BS can schedule arbitrary users for transmission with the fixed beams. If the users report per-subband CQI feedback that takes these beam constellations into account, the base station can perform link-adaptation that is only affected by normal fading as if no beamforming was used in the system. The BS can even achieve a certain coordination gain because it can avoid scheduling a user on subbands where colliding beams create high interference and rather pick a favorable subband with higher CQI. On the downside, the frequency-domain multi-user diversity is reduced because a user can only be scheduled with its preferred precoder on a smaller part of the total frequency bandwidth. However, as discussed in Section 5.4.1 and tabulated in Table A.2 in the Appendix, the 16 available PMIs in the LTE codebook are not chosen equally often for rank-1 precoding. Half of the PMIs (9–15) occur rarely and the remaining eight are also not equally probable depending on the deployment scenario and BS antenna array type. According to the frequency bandwidth (e.g., 20 MHz with 100 PRBs or 10 MHz with 50 PRBs), we can thus expect to have about 5–20 PRBs available per beam and subband if we exclude the rarely-used PMIs. The subbands for a specific PMI can also be distributed over the whole band in order to obtain as much diversity as possible.

Similar approaches have been proposed in the literature. With *opportunistic beamforming* (see [VTL02] and Section 5.4.1) also non-channel-adaptive beamforming is coupled with SINR feedback but without the multi-cell link adaptation aspect. And, for example, in [HvR09; vRH09] the authors suggest to independently cycle in the time domain through a fixed or only slowly varying set of beamformers in all cells. That way, a similar interference predictability is achieved with more diversity. However, a time-domain approach has the drawback that it adds a scheduling delay as the base station

has to wait for the next cycle until a known interference constellation occurs again. Note that while our proposed scheme is very simple, it is not fully compatible with the current LTE standard because a BS cannot signal which precoder shall be assumed on which subband for CQI computation. In LTE it is only possible to restrict the set of available PMIs to a subset that is then valid for the whole band.

In Section 6.2.3 we will show that already our simple scheme can improve the interference predictability and with it the spectral efficiency considerably. In the following section, we extend the beamforming to multiple users in one cell with MU-MIMO and also introduce a scheme that allows achieving a predictable interference situation.

6.2.2. Multi-User MIMO: Zero-Forcing and PU²RC

MU-MIMO is the extension of spatial multiplexing SU-MIMO techniques, which in our discussion so far were only considering a single BS–MS link, to multiple users. Instead of using N_{Tx} antennas for the transmission of up to N_{Tx} layers to one user (having $N_{\text{Rx}} \geq N_{\text{Tx}}$ Rx antennas), with MU-MIMO the BS transmits, for example, N_{Tx} individual single layers to N_{Tx} different users. A mixture of spatial multiplexing and MU-MIMO is possible as long as the total number of layers is not larger than the number of Tx antennas. In this thesis, we will only consider a single layer per user. This allows operating MU-MIMO also across highly correlated MIMO channels, which are common in outdoor deployments where the full spatial multiplexing gain cannot be realized (see Section 5.5.2). As another advantage compared to SU-MIMO, the MS could be equipped with fewer $N_{\text{Rx}} < N_{\text{Tx}}$ receive antennas. However, here we will assume $N_{\text{Rx}} = N_{\text{Tx}} = 4$ Rx antennas, which allow the MS to suppress the additional intra-cell interference resulting from transmission to co-scheduled users.

Even before MIMO spatial multiplexing was discussed in the literature [Tel99; FG98], the utilization of an antenna array at the base station for beamforming towards multiple users using the same frequency channel was known in the literature (see, e.g., [FN94]) under the name of *space division multiple access* (SDMA). In recent years, the term MU-MIMO has come to encompass SDMA and we will use it in the remainder. As mentioned in Section 2.1.3, *dirty paper coding*, which is the capacity-achieving MU-MIMO transmission scheme in the MIMO broadcast channel, is prohibitively complex to use in practice. Zero-forcing beamforming with perfect CSIT is a close-to-optimal alternative [YG06]. However, perfect CSIT is also infeasible, at least in FDD systems, due to the lack of channel reciprocity. Instead, we will thus only consider MU-MIMO schemes that can be implemented with a limited channel state feedback signaling overhead that is comparable to the PMI feedback in LTE TM4 and TM6.

One such scheme has been included as TM5 in the initial release of the LTE standard. Similar to rank-1 precoding with TM4 / TM6, the users feed back a PMI from the LTE codebook shown in Tables B.3 to B.7 to indicate the desired rank-1 beamformer, which is the first column of the selected matrix. The BS then selects two users (whose selected beamformers are as orthogonal as possible) and serves them simultaneously with their respective beamformers while distributing the Tx power equally among them. TM5

has several drawbacks [LYC+13]: First, the users compute their CQI and PMI feedback exactly as in TM6, that is, without taking the additional intra-cell interference from the co-scheduled user into account so that the BS has to perform some kind of CQI offsetting. Second, the MS is not informed about the PMI of the co-scheduled user. Since it cannot detect the intra-cell interference covariance based on measurements on the CRS, it is difficult to estimate the IPN covariance for IRC reception. Third, TM5 only allows for a wideband PMI feedback making it difficult to perform user grouping because variations in the frequency domain cannot be exploited. Fourth, it is not possible to dynamically switch between single user spatial multiplexing and MU-MIMO because that would involve switching between TM3 / TM4 and TM5, which can only be configured semi-statically by higher-layer signaling. As a consequence, TM5 was recently demoted to an optional feature [RP-121924; R2-126087] and future LTE deployments are expected to only support MU-MIMO in transmission modes TM8–TM10.

Motivated by the shortcomings of TM5, we examine two more advanced MU-MIMO variants. From a reference signal (CRS / DM-RS) and feedback signaling (CQI / PMI) perspective, the suggested methods comply with the existing LTE-Advanced standard. However, due to slightly different semantics regarding the feedback signaling and resource allocation, they cannot be applied directly without a minor modification to the standard.

Feedback-Based Zero-Forcing

With a zero-forcing MU-MIMO transmission scheme, the BS t' chooses the rank-1 precoders of co-scheduled users in a way that the precoding vector $\mathbf{p}_{k'}$ for a user k' is orthogonal to the channels $\mathbf{h}_{t'k}^T \in \mathbb{C}^{1 \times N_{\text{Tx}}}$ of all co-scheduled users $k \in \mathcal{D}_{t'} \setminus \{k'\}$ (see page 22 in Section 2.1.3 for the notation). Such precoders can be computed by stacking the channel vectors of all, for example, $N_{\text{MU}} = 4$ co-scheduled users as columns into a composite matrix $[\mathbf{h}_{t'k_1} \cdots \mathbf{h}_{t'k_4}]$, whose pseudoinverse yields the precoders \mathbf{p}_k^T as its rows, see e.g. [YG06]. Under the assumption of perfect CSIT, this means that the intra-cell interference between co-scheduled users is zero, hence the name. Here, we assume that the precoders are normalized to $\|\mathbf{p}_k\| = 1$ and that the Tx power is equally divided among the users.

An important aspect of ZF MU-MIMO scheduling is selecting suitable groups of co-scheduled users. Groups are chosen according to a desired metric, for example, the maximization of system spectral efficiency or some notion of proportional fairness extend to MU-MIMO [YG06; WPS+07]. As long as the channels from the BS t' to the group members are not already mutually orthogonal (which is highly unlikely for a limited number of users), the zero-forcing precoder $\mathbf{p}_{k'}$ cannot perfectly align with the serving channel $\mathbf{h}_{t'k'}^T$. The effective channel $\mathbf{h}_{t'k'}^T \mathbf{p}_{k'}$ thus has a reduced gain and hence lower SNR than with a normal beamformer that is solely matched to the channel $\mathbf{h}_{t'k'}^T$. Therefore, the SNR and with it the potential rate of each individual user in a group depend on exactly which other users are co-scheduled in the group. This makes the grouping process a difficult optimization problem that is often tackled by means

of heuristics like a greedy scheduler [DS05; Ell07b]. Also, for each potential group evaluated during the grouping process, a pseudo-inverse has to be computed.

If limited-rate feedback from the MS to the BS has to be used for CSI provisioning, the chosen zero-forcing precoders will not be perfectly orthogonal to the actual channels of co-scheduled users so that some residual intra-cell interference remains. Finding a feedback scheme with a good trade-off between low overhead and sufficient channel state accuracy has received a considerable amount of attention in the literature. In the introduction above, we implicitly assumed MISO vector channels $\mathbf{h}_{t'k'}^T \in \mathbb{C}^{1 \times N_{Tx}}$ where the MS only has a single antenna. In a typical LTE setting, however, we have MIMO channels $\mathbf{H}_{t'k'} \in \mathbb{C}^{N_{Rx} \times N_{Tx}}$ from the BS to the N_{Rx} antennas at the MS. Instead of feeding back the whole MIMO matrix whose size would depend on the actual number of Rx antennas, we can make the following simplification. For the reception of a single layer, a linear receiver at the MS employs a receive filter $\mathbf{w}_{k'}^H \in \mathbb{C}^{1 \times N_{Rx}}$ that is tailored to the serving channel and the interference conditions. Assuming the MS can estimate a suitable Rx filter $\mathbf{w}_{k'}^H$ in advance, the BS only needs information about the effective channel vector $\mathbf{w}_{k'}^H \mathbf{H}_{t'k'}$, which reduces the feedback requirements to that of a MISO vector channel [Jin06].

A classical way to feed back this effective channel, is using a pre-defined codebook \mathcal{C} of unit-norm channel direction vectors $\mathbf{u} \in \mathcal{C}$ together with standard CQI feedback. Different codebook choices are possible. Besides using a random codebook, a frequent approach is to use a *Grassmannian* codebook [LHS03], which consists of vectors that are chosen so that the minimum distance between any two unit-norm vectors is maximized. In addition to a codebook, also a suitable distance metric for choosing the codebook entry is required. The straightforward approach is to choose the codebook entry $\hat{\mathbf{u}} \in \mathcal{C}$ with the smallest angle to the effective channel $\mathbf{w}_{k'}^H \mathbf{H}_{t'k'}$. This minimizes the channel vector quantization error, which makes sense in a high-SINR regime since under such conditions the quantization error limits the SINR [DLU09]. However, as introduced by Trivellato et al. [TBH08a], maximizing the expected SINR, which is directly related to the feasible rate of the link, is the more relevant metric [DLU09; KDA+10]. It depends on the chosen precoder (which in turn depends on the feedback) as well as on the receive filter, which has to take into account the precoded serving channel in addition to the observed interference covariance.

For our MU-MIMO schedulers, we adopt the *maximum estimated SINR combiner* presented in [TBH08a]. According to that, the SINR estimation in (6.8) gives a lower bound on the expected SINR for user k' as

$$\text{SINR}_{k'}(\mathbf{u}) \geq \frac{\frac{P_{Tx}}{N_{Tx}} |\mathbf{w}_{k'}^H \mathbf{H}_{t'k'} \mathbf{u}|^2}{\mathbf{w}_{k'}^H \left(\frac{P_{Tx}}{N_{Tx}} \mathbf{H}_{t'k'} (\mathbf{I}_{N_{Rx}} - \mathbf{u} \mathbf{u}^H) \mathbf{H}_{t'k'}^H + \mathbf{R}_{IPN,k'} \right) \mathbf{w}_{k'}}, \quad (6.8)$$

where \mathbf{u} is a unit-norm codebook entry for channel vector quantization and $\mathbf{w}_{k'}^H$ is the receive filter that in turn depends on \mathbf{u} and on the inter-cell interference covariance $\mathbf{R}_{IPN,k'}$. As shown in [TBH08a], the denominator represents the filtered intra- and inter-cell interference assuming the maximum number of N_{Tx} users (reporting at least almost orthogonal channels) is co-scheduled with equal power $\frac{P_{Tx}}{N_{Tx}}$ per user.

After an exhaustive search over all possible codebook entries \mathbf{u} and corresponding filters $\mathbf{w}_{k'}^H$, the entry $\hat{\mathbf{u}}_{k'}$ that maximizes (6.8) is signaled to the BS:

$$\hat{\mathbf{u}}_{k'} = \arg \max_{\mathbf{u} \in \mathcal{C}} \text{SINR}_{k'}(\mathbf{u}). \quad (6.9)$$

In Section 6.2.3 we will present IMT_{Aphy} system-level simulation performance evaluation results for a ZF implementation using the above feedback strategy. Further implementation details regarding the employed codebook and greedy grouping algorithm will also be provided in Section 6.2.3.

But even before looking at the simulation results in detail, we want to highlight several drawbacks of MU-MIMO using zero-forcing beamforming. For example, the SINR estimation for feedback purposes is a difficult task for the mobiles contributing to bad link adaptation performance: Neither do the MSs know the actual residual intra-cell interference in advance, nor the fluctuating inter-cell interference. Similar to the situation with rank-1 precoding discussed in Section 5.4.3, zero-forcing beams in neighboring cells create a flashlight effect further worsening the link adaptation performance. The complexity of the group scheduling process has already been alluded to (note, however, that due to the finite feedback signaling choices, all possible groups and precoders could be pre-computed by the base station). Finally, the in principle arbitrary ZF precoders require additional demodulation reference signal overhead in the downlink, which lowers the spectral efficiency.

In the following, we will present a MU-MIMO variant that can overcome these limitations.

Multi-User MIMO with PU^2RC

In this subsection, we examine a MU-MIMO scheme that has been introduced by Samsung under the name *per unitary basis stream user and rate control* or later *per-user unitary rate control* (PU^2RC) in the early phase of the LTE standardization process [R1-030354; R1-050889; KKLK05]. However, it was ultimately not adapted for LTE MU-MIMO operation. In our evaluation here, which was published in [Ell12b], we focus on system-level and inter-cell interference aspects that have not been in the primary focus of the literature on PU^2RC (see, e.g. [KPL06; DPV07; DB07; TBH08b; WSS10; dFS10]).

The idea of PU^2RC is to restrict the available precoders to columns from a pre-defined codebook of unitary matrices. In fact, the 4×4 precoder matrices from the LTE 4 Tx antenna codebook (see Tables B.3 to B.6) are all unitary matrices and we will use a subset of the LTE codebook for our PU^2RC implementation. On a given time and frequency resource, each user from the group of co-scheduled user is assigned a different column from the same unitary precoding matrix. Similar to the closed-loop MIMO operation in LTE TM4, the users know the codebook of available precoders and indicate to the BS which precoder to use. To this end, with PU^2RC the MS selects and feeds back a desired column with a *column indicator* (CI) in addition to the desired PMI.

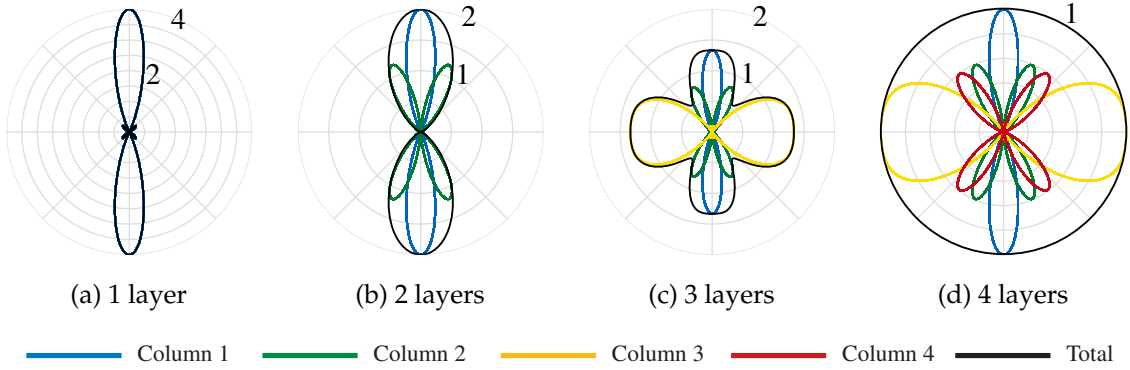


Figure 6.10.: Polar plot of the per-layer and total radiated power gain using 1–4 columns of the unitary PMI precoding matrix (see Table B.3) with a linear array of four $\lambda/2$ -spaced antennas

The feedback computation for PU²RC can rely on the exact same SINR expression given for ZF in equations (6.8) and (6.9). As a matter of fact, for a set of orthogonal co-scheduled feedback vectors (such as the columns of a unitary matrix), (6.8) holds with equality [TBH08a]. Note that the term $\frac{P_{\text{Tx}}}{N_{\text{Tx}}} \mathbf{H}_{t'k'} (\mathbf{I}_{N_{\text{Rx}}} - \mathbf{u}\mathbf{u}^H) \mathbf{H}_{t'k'}^H$ in the denominator of (6.8) is the intra-cell interference covariance from the other unitary matrix columns besides \mathbf{u} . If we add the serving precoder's transmit covariance $\mathbf{u}\mathbf{u}^H$, we obtain the identity transmit covariance we expect from PU²RC. Apart from the consideration of the intra-cell interferers, the general feedback computation is identical to the rank-1 precoder feedback, for example, in TM6. The best feedback PMI, column indicator (CI), and CQI can be found with a simple exhaustive search over the codebook.

The main advantages of PU²RC are the following:

1. **Precise SINR estimation:** The MS receiver not only knows its serving precoder but also the possible precoders of all co-scheduled users in advance. This allows an exact estimation of the (worst-case) intra-cell interference and hence overall SINR for CQI feedback and link adaptation purposes.
2. **Intra-cell interference suppression:** Due to the PU²RC conventions, the BS is forced to transmit to co-scheduled users using the remaining columns from the chosen unitary matrix. The MS thus implicitly knows the intra-cell interference covariance and can perform interference rejection combining to suppress it.
3. **Reduced or eliminated flashlight effect:** The characteristic property of a unitary matrix \mathbf{U} is that its covariance is the identity $\mathbf{I} = \mathbf{U}\mathbf{U}^H$. This means that the interference covariance $\mathbf{R}_{tk'} = \mathbf{H}_{tk'} \mathbf{U}_t (\mathbf{H}_{tk'} \mathbf{U}_t)^H = \mathbf{H}_{tk'} \mathbf{H}_{tk'}^H$ of BS t as seen by receiver k' is independent of \mathbf{U}_t and only depends on the channel $\mathbf{H}_{tk'}$. Thus, the inter-cell interference covariance \mathbf{R}_{IPN} does not depend on which PMI was scheduled on a certain physical resource (PRB and TTI). This allows much better interference covariance estimation (the channel's covariance $\mathbf{H}\mathbf{H}^H$ can be expected to be correlated for a region larger than a TTI and PRB, in most cases) further contributing to better link adaptation performance. If a BS cannot fill the group and schedules

$N_{\text{MU}} < N_{\text{Tx}}$ users, a reduced flashlight effect remains. As Figure 6.10 shows, the directivity of the total radiated power reduces as more columns of a unitary precoder are used until the total power is uniform when all columns are utilized.

4. **Low scheduling and grouping complexity:** In contrast to a ZF scheduler, the PU²RC grouping algorithm is straightforward because the rate of a group member does not depend on exactly which other users are co-scheduled since the possible co-scheduled precoders were already taken into account during the feedback computation process. Also, no matrix inversions are necessary and no CQI feedback correction is required as for ZF scheduling where the gain of the effective serving channel depends on the ZF precoders.
5. **Reduced reference signal overhead:** When arbitrary precoders are used as with ZF, each user needs demodulation reference signals (DM-RS) in addition to channel estimation reference signals (CSI-RS) to be able to demodulate the precoded transmission. With PU²RC, cell-specific reference signals (i.e., common to all users) together with codebook entry signaling can be used. As we will detail in Section 6.2.3, this leads to a lower overhead for PU²RC compared to ZF precoding in an LTE-Advanced system.

Besides the above advantages for PU²RC, also some drawbacks have to be mentioned. For the PU²RC scheduler to be able to find suitable groups with the maximum number of co-scheduled users, enough users with diverse feedback and a not-too-large codebook are needed. If the scheduler cannot fill a group, two general approaches are possible:

On the one hand, with a “fill” variant, the BS can force the group to be completed, and the precoding matrix to be unitary, by assigning the missing columns to other users who originally picked different PMIs or CIs. The downside is that users have to be scheduled with precoders for which no suitable CQI feedback is available.

On the other hand (“normal”), it can leave columns from the unitary matrix unused. This means that also the full BS Tx power is not exhausted because the power for the remaining group members cannot be increased due to the cell-specific demodulation reference symbols, which again causes a light flashlight effect. But as an advantage, the overall interference level is lower and interference-rejecting receivers can spend remaining degrees of freedom on suppressing strong inter-cell interferers. Also, the link adaptation using the PU²RC feedback works. However, the PU²RC CQI feedback will be too conservative, but this can be corrected by the OLLA mechanism.

6.2.3. Performance Evaluation of ZF and PU²RC MU-MIMO

In this section, we present the performance evaluation of the predictable beamforming (i.e., the TM6 variant with fixed PMIs introduced in Section 6.2.1), and of three flavors of MU-MIMO scheduling based on the *maximum estimated SINR combiner* [TBH08a] feedback defined above in equations (6.8) and (6.9): zero-forcing, “normal” PU²RC that co-schedules only as many users per resource as it has users with compatible feedback,

and a robust PU²RC variant that enforces unitary precoding matrices by filling groups also with users reporting different PMIs or CIs in order to reach N_{Tx} co-scheduled users. For comparison with rank-1 precoding and SU-MIMO with spatial multiplexing per user, we show the corresponding TM4 and TM6 results. Before that, we will briefly describe selected implementation details, discuss the overhead requirements of the different schemes, and give an overview of simulation parameters.

Selected Implementation Details and Overhead Requirements

For the predictable beamforming implementation, we use our TM6 implementation and only change the feedback computation. In each cell, a fixed cell-specific PMI-to-PRB mapping is randomly selected and the users report their frequency-selective CQI feedback taking this mapping into account. Each BS adheres to the cell-specific fixed PMI feedback per PRB, but can schedule arbitrary users onto the PRBs so that in each cell a constant transmit precoding per PRB results. Here, we distribute the 16 PMIs from the 4 Tx codebook randomly to the PRBs.

For both the PU²RC and zero-forcing scheme based on the maximum estimated SINR feedback, we use an identical codebook \mathcal{C} of unitary matrices that consists of the PMIs $\{0, 1, 4, 5\}$ from the standard LTE Rel-8 codebook for 4 Tx antennas shown in Tables B.3 and B.4. As an extension to the standard LTE signaling, we assume that besides a 2-bit PMI index, also the 2-bit column indicator is fed back. To address the 16 possible (PMI, CI) feedback vector combinations, we thus need the same overhead as in TM4 and TM6 for the 16 PMIs making for a fair comparison. In addition to a frequency-selective feedback of preferred PMI and CI indices, the users report CQIs to their base stations based on estimated SINRs derived using equations (6.8) and (6.9).

The MU-MIMO scheduling in the BS is performed using a straight-forward extension of the proportional fair scheduling metric (5.1) (see page 111) to groups of co-scheduled users. On a PRB f_{id_x} at TTI n , the BS t' selects that set $\mathcal{D}_{t'}^*(f_{\text{id}_x}, n)$ of users to be co-scheduled that maximizes the sum of the PF metrics of each group member:

$$\mathcal{D}_{t'}^*(f_{\text{id}_x}, n) = \arg \max_{\mathcal{D}_{t'}} \sum_{k \in \mathcal{D}_{t'}} \frac{r(k, \mathcal{D}_{t'}, f_{\text{id}_x}, n)}{\overline{\gamma_{k,n}}^\beta}, \quad (6.10)$$

where the throughput history $\overline{\gamma_{k,n}}^\beta$ is computed as for the single user case. Regarding the expected rate $r(k, \mathcal{D}_{t'}, f_{\text{id}_x}, n)$ of a user k , the ZF scheduler has to take into account that the precoder \mathbf{p}_k and thus the rate depend on the set of co-scheduled users $\mathcal{D}_{t'}$. As a consequence, with zero-forcing scheduling the optimal group size $N_{\text{MU}}^* = |\mathcal{D}_{t'}^*|$ can vary. The best grouping can be found by an exhaustive search over all possible groupings $\mathcal{D}_{t'}$ where the number of co-scheduled users varies from one to N_{Tx} . Due to the huge number of possible groupings [Ell07b] and the need to perform matrix inversions, the optimal ZF scheduler is computationally complex. The implemented ZF scheduler thus performs a greedy grouping that, starting with an empty group, always adds that user from the set of remaining users that yields the largest MU-MIMO PF scheduling metric increase. In contrast, if the PU²RC scheduler assigns its fed back

PMI and column to a user, the expected user rate $r(\text{CQI}(k, f_{\text{idX}}, n))$ only depends on the fed back CQI value, but not on the identity of the co-scheduled users. This allows a very simple grouping algorithm that groups those users who report the same PMI but different CIs. If multiple user report the same (PMI, CI) combination, the user having the best ratio of expected rate (i.e., the highest CQI) and history exponent γ can be picked independent of the co-scheduled users.

The PU²RC schemes rely on the same cell-specific reference symbols (CRS) that are used for TM4 and TM6 as well. The receiver at the MS estimates the serving channel $\mathbf{H}_{\ell'k'}$ and receives the PMI of the used precoder $\mathbf{P}_{k'}$ via downlink signaling so that it can compute the effective serving channel $\mathbf{H}_{\ell'k'}\mathbf{P}_{k'}$. As mentioned above, the ZF scheme employs arbitrary ZF precoders \mathbf{p}_{ZF} so that it requires user-specific demodulation reference symbols to allow the MS to estimate the effective serving channel $\mathbf{H}_{\ell'k'}\mathbf{p}_{\text{ZF}}$. This has two consequences for our modeling here:

On the one hand, we use different reference signal overhead assumptions (see also Section 2.3.2 and Fig. 2.12). The CRS-based schemes are modeled as usual with CRS overhead for four antenna ports. For the ZF scheme, we use the Rel-10 DM-RS. One set of DM-RS (e.g., the blue-colored one in Fig. 2.12b) consisting of 12 REs is sufficient for our 4 Tx antenna ZF scheme. In order to reduce the overhead and as shown in Fig. 2.12b, we only model CRS for 2 antenna ports. But, compared to CRS with 4 Tx ports and neglecting that some CSI-RS would also be needed for channel estimation for CSI feedback purposes, this still amounts to a net increase of eight occupied REs outside the control region within one TTI and PRB. Zero-forcing thus has an almost 7% higher PHY overhead than the CRS-based schemes.

On the other hand, the intra- and inter-cell interference covariance under the ZF scheme is much harder to estimate than in the case of PU²RC. We account for this disadvantage by modeling ZF with an estimation error based on the Wishart model introduced in Section 4.3.2 assuming $N_{\text{sample}} = 12$ available samples within one PRB and TTI (i.e., one sample per DM-RS RE). In contrast, the PU²RC scheme is modeled assuming perfect interference covariance estimation.

Further implementation details, such as the handling of HARQ retransmissions, can be found in the published source code [IMTAphyWeb; IMTAphyLP]. Table 6.4 summarizes the simulation parameters for the results presented in the following.

Simulation Results

For our simulations we have chosen the UMi scenario due to the low user mobility, which allows channel-adaptive scheduling as demonstrated in Chapter 5. First, we examine the distribution of normalized throughput, average SINR, and an idealized link adaptation SINR mismatch per user in a scenario with on average 25 users per cell. After that, we demonstrate how the grouping performance of the MU-MIMO schedulers varies with different numbers of users per cell.

Figure 6.11 shows the distribution of the per-layer average SINR per-user. As expected, the single-user beamforming variants achieve the best SINRs since they can exploit

Table 6.4.: Simulation parameters for the simulations presented in Section 6.2.3

Parameter	Value
Scenario	UMi, users with default speed of 3 km/h
Number of drops	5 for $N_{MS} > 30$, 10 for $N_{MS} > 15$, 20 otherwise
Number of MS	on average $N_{MS} \in \{5, 10, 15, 20, 25, 30, 40, 50\}$ users per cell
BS and MS antennas	4×4 linear arrays with $\lambda/2$ spacing on both sides (BS config C)
Receiver type	MMSE-IRC
Transmission schemes	TM4 SU-MIMO; standard TM6 beamforming and predictable beamforming; normal PU ² RC, “fill” PU ² RC, and ZF based on maximum estimated SINR feedback [TBH08a]
Link adaptation	with outer-loop link adaptation (OLLA)
Subband sizes	2 PRBs per subband for CQI and PMI feedback
Scheduler	proportional fair with $\beta = 1$, adapted to MU-MIMO, see (6.10)
HARQ	Chase combining, asynchronous and adaptive
Channel estimation	perfect serving channel estimation interference covariance estimation perfect or based on Wishart model with 12 samples, see Section 4.3.2
Codebook	LTE Rel-8 codebook for 4 Tx ports with 16 PMIs (Tables B.3 and B.4) and subset \mathcal{C} thereof consisting of PMIs $\{0, 1, 4, 5\}$
RS overhead	CRS for 2 Tx ports and 12 REs DM-RS per PRB and TTI for zero-forcing; CRS for 4 Tx ports otherwise




Figure 6.11.: Distribution of SINR per user and layer for different MU-MIMO and reference schemes in a UMi scenario with on average 25 users per cell




Figure 6.12.: Distribution of SINR estimation mismatch for different MU-MIMO and reference schemes in a UMi scenario with on average 25 users per cell

the array gain and restrict interference to certain spatial directions. The next best SINR curve belongs to single-user spatial multiplexing with TM4, which can adaptively use higher transmission ranks, but in practice is mostly limited to 1–2 layers (see Fig. A.7b in the Appendix). The SINRs for all MU-MIMO schemes are significantly worse due to the additional intra-cell interference and the power sharing between users. Nevertheless, the MU-MIMO SINR distributions span the useful range of SINRs, whereas the single-user beamforming SINRs often reach the saturation level of the available LTE modulation and coding schemes. As explained above, the ZF scheme should be modeled with imperfect interference covariance estimation. This has a significant impact on the obtained SINRs due to the lower interference estimation accuracy with the Wishart model compared to a perfect estimation, as shown in Fig. 6.11. An even bigger impact on SINRs stems from the average MU-MIMO group sizes because the more users are grouped together, the less power per user is available, and the higher the intra-cell interference. Here, the normal PU²RC scheme yields the best SINR distribution because it can find on average fewer compatible users than the greedy ZF scheduler adds users to the groups. The worst SINR distribution is obtained with the robust PU²RC variant since it fills groups to have $N_{\text{MU}} = N_{\text{Tx}}$ users by adding incompatible users.

In contrast to its bad average SINRs, the robust group-filling PU²RC scheme achieves the best, that is, smallest SINR estimation mismatch, which is shown in an idealized way (before any OLLA correction) in Fig. 6.12. Here, we plot the *probability density function* (PDF) of the actual received transmission SINR minus the unquantized SINR estimate (note that in the simulation the link adaptation is performed on the normal quantized feedback values). The remaining small variations for “fill” PU²RC and predictable beamforming are due to the fast-fading channel and would disappear completely for a static channel as shown in [Ell12b]. The other schemes exhibit a higher




Figure 6.13.: Distribution of normalized user throughputs for different MU-MIMO and reference schemes in a UMi scenario with on average 25 users per cell

variance due to the fluctuating interference (flashlight effect). The impact of interference covariance estimation errors is also clearly visible comparing the mismatch distributions of the ZF scheme with perfect and Wishart-based estimation accuracies.

Applying the proportional fair channel-adaptive scheduling to groups of users alleviates the single-user tendency to overestimate the SINR when trying to “ride on the peaks” of the fading channel. Thus, the absolute average SINR mismatch, and with it the need for OLLA compensation, is significantly smaller with the MU-MIMO schemes than in the case of TM4 and TM6, as shown in Fig. 6.12. The MU-MIMO transmission schemes also exhibit a smaller variance of the SINR mismatch because transmitting multiple layers in parallel makes the radiated interference more uniform as motivated by Fig. 6.10. Since the PU²RC feedback assumes the maximum number of co-scheduled users, the SINR mismatch for the normal PU²RC scheme is positive because usually fewer interfering users are co-scheduled.

Finally, we piece the above results together and look at the resulting throughputs depicted in Fig. 6.13 and listed in Table 6.5. The zero-forcing reference result, which only considers CRS overhead and assumes perfect interference covariance estimation, performs best. However, already when adding the DM-RS overhead, the performance falls behind that of the normal PU²RC scheme. The situation worsens as we consider the Wishart error model with $N_{\text{Sample}} = 12$, but the realistically modeled ZF scheme still achieves a slightly better result than the “fill” PU²RC variant. With on average 25 users per cell, the latter cannot benefit enough from the better link adaptation performance to offset the lower SINRs we observed above. Since the applied ZF grouping heuristic is a simple greedy one, better results could be achieved with a more advanced strategy at the cost of a higher complexity. Overall, Table 6.5 shows a significant advan-

Table 6.5.: Cell spectral efficiency and cell-edge user throughput in bit/s/Hz in the UMi scenario for MU-MIMO and reference schemes with 25 users per cell

	TM6	TM4	TM6 pre- dictable BF	Zero-forcing DM-RS, 12 samples	Fill PU ² RC	Normal PU ² RC
Cell spectral efficiency	2.76	2.80	3.15	3.40	3.27	4.00
Cell-edge user throughput	0.046	0.038	0.055	0.051	0.045	0.060

tage of the MU-MIMO transmission schemes over single-user MIMO with either TM4 or TM6 assuming equal feedback overhead: The normal PU²RC scheme achieves a 43–45 % gain in cell spectral efficiency and a 30–58 % gain in cell-edge user throughput compared to TM4 and TM6 when on average 25 users are available per cell.

With respect to the simple predictable beamforming scheme introduced in Section 6.2.1, we observe a significant throughput gain over the conventional LTE rank-1 precoding scheme (TM6), which is solely due to the better link adaptation. More sophisticated implementations (e.g., semi-statically adapting the used PMIs to the channel conditions of the users as suggested in [Ell10]) could probably improve the performance further. However, considering the significant gap also of the predictable rank-1 beamforming to the presented MU-MIMO schemes, it is evident that, at least in the outdoor scenarios suffering from higher channel correlation, MU-MIMO transmissions play a crucial role for any prospective future spectral efficiency increases.

The advantage of MU-MIMO over SU-MIMO is also visible in the final two plots in Fig. 6.14 that show the average SINR and the idealized SINR estimation mismatch as well as the cell spectral efficiency over a varying number of users per cell. The MU-MIMO schemes outperform the SU-MIMO ones already for small number of users, but as the average number of users per cell grows towards the maximum of 50 users per cell in the plots, the MU-MIMO multiplexing gain increases. With larger groups, the average SINRs decrease (except for the fill scheme, which actually gains because it has to force less incompatible users) and the SINR estimation mismatch is further reduced. Also the predictable beamforming variant benefits considerably because more users allow for more multi-user diversity which is otherwise restricted by the fixed PMI-to-PRB mapping. Note that in the plot the TM4 performance decreases with growing number of users per cell, which is due to the very aggressive PF scheduler that leads to increasing SINR mismatches.

Given the clear advantages of neighbor-considerate transmission schemes (such as predictable beamforming and PU²RC) and of MU-MIMO in general, the question remains why these or similar schemes are not yet included in the current LTE-Advanced standard or deployed in real-world networks. The most likely answer is that current cellular systems can still be operated using simpler methods. For example, many of the spectral-efficiency-increasing techniques introduced in Chapter 5 (such as channel-

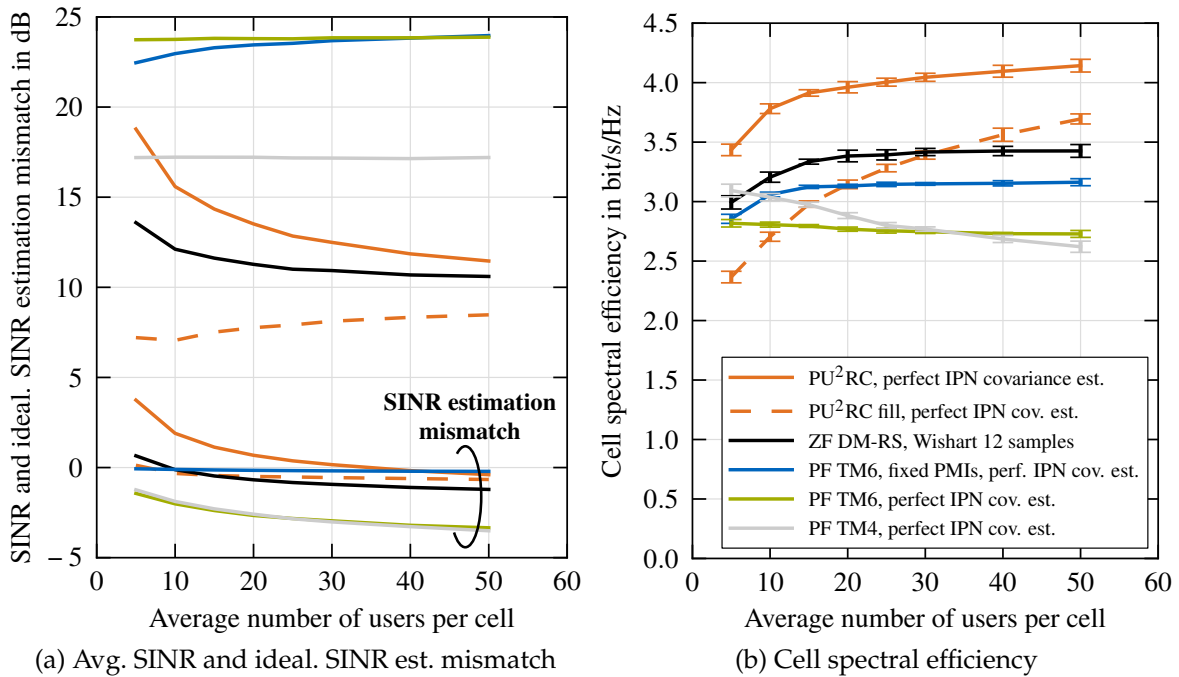


Figure 6.14.: Impact of multi-user diversity on average SINRs, idealized SINR estimation mismatches, and consequently throughput with presented MU-MIMO schemes with TM4 and TM6 as a reference

dependent frequency-selective scheduling or closed-loop precoding with TM4 or TM6) are not yet widely observed in currently deployed networks; presumably because they are more complex and not yet needed. The traffic load in current LTE networks is mostly too low and, in addition, the traffic pattern is too bursty. The benefits of the above schemes, however, are mainly applicable in cells that are fully loaded because otherwise the scheduling activity (i.e., whether a BS transmits at all) dominates the interference fluctuations and the spatial structure of the interference is only of secondary concern. After standardization, one intermediate step to ease the introduction could thus be to combine frequency reuse coordination techniques discussed in Section 6.1 with the neighbor-considerate schemes introduced above. By bundling the scheduling activity of an underutilized cell in one part of the bandwidth, no inter-cell interference is created in the other part of the bandwidth, whereas the used part exhibits a predictable interference pattern.

What is further important to note here, is that the presented performance evaluation is another example underlining the importance of a careful and detailed system-level modeling: In a simplistic evaluation setup (or when evaluating even only a single isolated cell), zero-forcing would outperform the normal PU^2RC scheme. Only if we model and consider the link adaptation performance in the face of fluctuating intra- and inter-cell interference, and furthermore take reference signal overheads and estimation accuracies into account, we are able to conclude that PU^2RC is the better choice in a real-world setup.

6.3. Outlook to Future Improvements in LTE-Advanced and 5G Systems

At the beginning of the new millennium, Martin Cooper established *Cooper's Law* stating that, in the first 104 years since Guglielmo Marconi's first radio transmission in 1895, the area spectral efficiency has doubled every 30 months [Coo]. He argued that in the past 50 years the corresponding growth by a factor of more than one million was mostly due to smaller cell sizes (1600-fold) and more available spectrum (25-fold). Frequency division operation and improved PHY layer techniques, which did until then not include gains due to MIMO operation, contributed only a factor of five each [Coo; DHL+11].

In the end, only history will tell if the growth in wireless as predicted by Cooper's Law will continue in the future and how it will be achieved. In time for the upcoming WRC in 2015, the ITU-R is currently putting together an "IMT.VISION Framework and overall objectives of the future development of IMT for 2020 and beyond" document where it solicits contributions from external organizations [Wor13]. Initial visions have already been expressed both by industry and academia. For example, Nokia Siemens Networks [Nok11] envisions a 1000-fold increase of area spectral efficiency until 2020, where a 10-fold improvement should come each from the areas of spectrum availability, system spectral efficiency, and cell density. Other sources [ZM13] come to similar conclusions regarding the total growth, but differ regarding the relative importance of the categories. In the following, we will briefly review a few promising directions in which the data rates in cellular networks can be improved for future versions of LTE-Advanced and its 5G successors.

6.3.1. Deployment and Spectrum Aspects

Having more and thus smaller cells will also be an effective way to achieve higher data rate per user in the future. However, adding more sites, base station equipment, and the necessary backhaul is very expensive for operators. Help could come from femtocell deployments in a heterogeneous network as introduced in Section 6.1.3. A different kind of HetNet deployment is currently studied [36.932] as a major component of LTE Rel-12. The so-called *small cells* are operator-deployed low-power cells that are supposed to complement an existing macro cell network. Since they are intended to serve only a small area with high traffic demand, they need to provide only a reduced range and can thus make use of higher frequency spectrum (e.g., in the 3.5 GHz band). High frequencies are in general not preferred for wide area coverage due to high pathloss at these frequencies. Nevertheless, there is currently increased research activity [RSM+13] to make even higher frequencies up to 60 GHz available for cellular systems because a large amount of potential spectrum exists in these "mm-wave" bands.

Another approach that would benefit from high frequencies and thus smaller antennas is *massive MIMO* [Mar10; HCPR12] where a huge number of BS antennas would allow

to serve many users in parallel in a MU-MIMO fashion. In the same direction, but closer to practical deployment, are studies [RP-122034] performed within 3GPP on 3D MIMO. The main idea there is to have 2D antenna arrays at the BS that allow using the vertical dimension for elevation-based beamforming as well as vertical cell sectorization. Both approaches enable the BS to serve more users in parallel in the same area, either through higher degrees of MU-MIMO or by creating more cells.

The easiest way to increase cellular data rates is, of course, the provision of more spectrum by the ITU and national regulators. Besides the 700 MHz band, which is already slated for worldwide availability, upcoming WRCs could make more (conventional) spectrum available for cellular communications. Another approach seeks to exploit already otherwise allocated but underutilized spectrum. To use these *white spaces* in the radio spectrum, *dynamic spectrum access* [ZS07] methods have been devised, which were often tailored to dynamic and opportunistic short-term spectrum usage by non-cellular devices. Recently, *licensed shared access* [KBDG12] has been proposed as a way to allow incumbent spectrum holders to sublicense their spectrum in a controlled way to cellular operators, who could then use it for cellular communication systems like LTE-Advanced.

6.3.2. Link Efficiency Aspects

On the link level, the potential future gains are limited in line with the development in the past. Clearly, massive MIMO can improve the link efficiency. Not so much by enabling more spatial layers but rather by enabling higher SINRs at the MS due to the beamforming gain at the BS. In order to benefit from higher SINRs, the introduction of higher-order modulation schemes, namely 256-QAM, would be a natural extension to LTE-Advanced. Another area of improvement is reducing the overhead in LTE, which accounts for a significant loss in spectral efficiency as discussed in Sections 2.3.2 and 5.5.2. In that regard, an LTE Rel-12 work item [RP-122028] aims at establishing a so-called *new carrier type*. Such a *lean carrier* [HLKC13] could solely rely on user-specific DM-RS and on the CSI-RS scheme introduced in Rel-10, see Section 2.3.2. This would allow to remove the CRS overhead. In addition, the legacy PDCCH, which is transmitted over the whole bandwidth in the first 1-4 out of 14 symbols per TTI, could be replaced by the more efficient *enhanced physical downlink control channel* (ePDCCH) introduced in Rel-11 [YWW13]. Such a new carrier type would not support legacy mobiles, but it could be the ideal choice for the above-mentioned small cells, providing additional localized high data rate coverage in addition to legacy macro cells. If a new carrier design is primarily tailored to small cells, the cyclic prefix length could be reduced [MPR+12] because expected delay spreads in small cells are shorter than in regular macro cells. That way, a few percent of overhead could be saved.

A link performance improvement that does not require standardization support is the use of near-ML-detection receivers such as sphere decoders. In addition, advanced receivers can perform interference cancellation on the cell-specific reference signals (CRS) from other known cells that interfere with a desired PDSCH transmission.

Much further out on the horizon is the use of full-duplex radio transceivers, prototypes of which have been presented recently [JCK+11]. The implications of full-duplex operation are not only that the same spectrum can be used simultaneously for both uplink and downlink transmission, achieving an up to twofold spectral efficiency gain; it would also allow the provision of nearly instantaneous CSIT at the base station because the MS can send uplink pilot signals on the same frequency band that is used in the downlink. This could allow the BS to estimate the DL channel based on channel reciprocity. Compared to FDD operation, this avoids the problem of limited and delayed CSI feedback. Compared to TDD operation, where channel reciprocity could also be used in the DL phase after estimating the channel in a previous UL phase, this avoids CSI aging effects resulting from the switching delays between the DL and UL transmission phases. On a side note, full-duplex radios could also enable new MAC schemes in WLANs because the ability to send and receive at the same time can help to eliminate the *hidden terminal* problem. With full-duplex radios, the receiving WLAN station can alert its surroundings by sending channel state feedback or a simple busy tone [JCK+11].

6.3.3. System-Level RRM Efficiency Aspects

Higher gains in spectral efficiency have to come from system-level improvements. Especially rates at the cell edge could be improved to achieve a more homogeneous distribution of feasible rates over the whole cell area. HetNet deployments, which add further small low-power base stations inside a regular macro cell, promise to improve rates in high-traffic or cell-edge areas. An early proposal in this direction was the use of relays [PWS+04]. In contrast to femtocells, they do not need a wired backhaul connection because the backhaul is fed over the air either in-band or out-of-band from the macro base station. As the feeder link also needs spectrum, the net spectral efficiency gain (or loss) of relay-enhanced cells depends on the quality of the feeder link, careful planning of the relay deployment, and the number of relays per cell [HCM+12].

Coordinated Multi-Point Transmission (CoMP)

On a system level, the biggest potential for spectral efficiency improvements in multi-cell systems lies in the ability to remove inter-cell interference while maintaining a full frequency reuse. This can be achieved if the base stations in neighboring cells are not operated independently, but are rather operated jointly as one large distributed antenna array serving a single super-cell, which is sometimes called a *cluster*. Such schemes have found large interest in the research community [MF11] under names like *network MIMO* or *network coordination*. Recently, they have made their way into standards such as LTE Rel-11 under the name CoMP. Initially, there has been a huge hype around CoMP as early publications [CS03; KJV06] showed that coordinated network MIMO could offer, for example, a five-fold capacity increase in the downlink

over uncoordinated MIMO transmissions. These studies assumed that perfect multi-cell CSIT is available at all *transmission points* (TPs) so that users could be served in a MU-MIMO way with ZF and DPC processing at the transmitter. Clearly, perfect multi-cell CSIT will never be available in an actual system as this is already infeasible in single-cell MIMO contexts. In the CoMP setting, in addition, the CSI for links to all geographically separated TPs would have to be estimated and made available at all TPs in a timely manner. Also, the payload data would have to be provided to all TPs in order to enable a jointly processed transmission from all antennas. These high demands for low-delay backhaul capacity [MF11, Ch. 12] are one of the main reasons why the initially promising gains could not yet be implemented in standards even though extensive field trials showed promising results too [IDM+11]. Additional real-life challenges [YHHM13] are multi-cell CSI acquisition with the corresponding overhead for training and feedback on the air interface, as well as other issues such as per-BS power constraints and asynchronous interference.

In order to relax these requirements, the standardization of CoMP features in LTE Rel-11 [LKL+12] introduced two less demanding variants in addition to the full network MIMO approach, which is called *joint transmission* (JT) in 3GPP parlance. Both of them do not require to exchange user data with all participating TPs and have significantly reduced CSI requirements. With *coordinated scheduling / coordinated beamforming* (CS/CB) neighboring cells can coordinate their beamforming in order to avoid conflicts. With *dynamic point selection* (DPS) a mobile can be switched from one serving TP to the other on a millisecond timescale without having to perform handovers. In addition, a mechanism called *dynamic point blanking* (DPB) allows to dynamically mute a non-serving transmission point to remove its interference. To facilitate these new features in Rel-11, the CSI feedback mechanism has been extended. It is now possible to configure an MS to provide multiple CSI (CQI, PMI, RI) reports, which are measured on separate REs (*measurement sets*). With an appropriate configuration of CSI-RS transmission (and muting) between the antennas of different TPs on these different measurement sets, the network can obtain CSI feedback that allows JT, CS/CB, or DPS CoMP schemes [LKL+12].

Instead of providing a high-capacity backhaul between separate base stations, CoMP deployments are likely to use *distributed antenna systems* [HPWZ13] where so-called *remote radio heads* (RRHs) are connected via fiber to a central base station or *cloud RAN* [LWKL11]. The RRHs are then only antennas with an amplifier and all multi-cell physical-layer processing is done at one place.

The performance of CoMP in LTE has been investigated in homogeneous scenarios with intra-site as well as multi-site coordination. In addition, heterogeneous scenarios have been considered, where low-power base stations are coordinated. The evaluation results obtained by different 3GPP members during the study phase are summarized in [36.819]. Compared to the considerable complexity of CoMP, they show mostly modest gains in the downlink and differ widely between companies. Overall, no convincing conclusion has been presented for the downlink performance. In the uplink, however, joint reception is agreed to offer a “considerable gain” over uncoordinated reception. A further advantage of uplink CoMP with joint reception is that, for example, no feedback

signaling on the air interface needs to be standardized. Independent groups have also presented performance evaluation results of LTE Rel-11 CoMP [SGP+13; YHHM13]. The authors in [SGP+13] show cell-average and cell-edge gains of 15 % for JT with intra-site CoMP and gains of 18 % for multi-site JT. Results for heterogeneous scenarios show a single-digit gain in average throughput. For the cell-edge throughput in HetNets, a 19 % or 28 % gain is achieved if the coordination is performed for the low-power nodes within one sector or across three sectors, respectively. Yang et al. [YHHM13] show system-level simulation results in a UMa scenario with three coordinated cells per cluster. According to their results, cell center users do not benefit from JT CoMP when overheads are considered, and also the cell-edge users only realize a modest gain from JT CoMP. As a conclusion, we can summarize that with LTE Rel-11, the gains from CoMP in the downlink are still far from the theoretical bounds. Due to the long list of practical challenges [YHHM13], it is clear that the theoretical benefits will not be reached completely. Whether it will be possible to get closer to them, will be an interesting topic for future research.

A signal processing concept that received a considerable amount of attention in the past years is *interference alignment* [CJ08]. Applied to the downlink of cellular networks [TG09; SHT11], it can be seen as a coordinated multi-point transmission approach that performs a linear precoding at each BS in a way that for each MS receiver the interference from all undesired BS transmitters aligns in a lower-dimensional subspace. The mobiles can then use simple linear ZF receive filters to suppress all the interference while receiving the desired signal component that is orthogonal to the interference subspace without interference from the aligned undesired transmissions. Similar to JT strategies, full CSIT of all channels is required, but the data is only needed at the serving BS because each user is served separately. Given the high CSIT requirements, currently limited suitability at the low SNRs predominant in cellular networks, and other issues [EPH13], it remains to be seen if it can be successfully employed in cellular contexts.

Potential System-Level Improvements with LTE Rel-12 and Beyond

A few topics currently under investigation and standardization for LTE Rel-12 and beyond also promise to enhance the spectral efficiency on a system level. As already mentioned, small cells are a focus of Rel-12 standardization. From a system-level perspective, small cells that are deployed within the coverage of macro cells allow an interesting hybrid *dual connectivity* operation mode. In such a configuration, a mobile always stays connected to the macro cell and exchanges control messages with the macro cell. But high data rate transmissions are performed via the small cells, which are called *booster* or *phantom cells* [NNB+13] in this context. On the one hand, a small coverage area with few competing users, a lean carrier, and additional spectrum allow high data rates in the booster cell. On the other hand, the dual connectivity to the macro cell allows a mobile user to be *always on* without requiring frequent handovers between small cells when no high data rate connection is ongoing. Also, the lean carrier in the small cell does not need to sacrifice resources for control signaling.

Another approach to offload traffic from macro cells is the ability to perform *device-to-device* (D2D) transmissions directly between two geographically close mobiles. The degree to which spectral efficiency can be improved by this depends significantly on the amount of local traffic that could be carried by one D2D link instead of occupying separate UL and DL cellular links. The 3GPP is currently investigating this under the name proximity services (ProSe) [23.703] after a feasibility study [22.803] has already been concluded. There are several aspects under discussion [FDM+12]. On the one hand, it has to be determined which spectrum is used. The D2D link could use the macro cell's licensed spectrum or resort to unlicensed spectrum employing WLAN protocols. On the other hand, the degree of network control has to be decided. When using licensed spectrum, the BS might want to have tight control over the resource allocation. In unlicensed spectrum, the role of the BS might be reduced to providing security and discovery services. The latter would nevertheless significantly help the devices to save energy as they would not need to constantly scan for other devices and send beacon signals themselves. Additionally, a big motivation for D2D comes from public safety use cases where D2D links would allow communications even if the infrastructure is not available, for example in disaster situations.

Furthermore, the system spectral efficiency could be improved by *network-assisted interference cancellation and suppression* (NAICS) support from the network. The current LTE control channel design makes downlink resource allocation information only available to the desired receiver because that receiver's temporary identifier is used to scramble the control information. If more information about the structure of the interference, for example the employed modulation scheme, would be available to an interfered user, the processing in the receiver could be improved because parts of the interference would not have to be considered as AWGN. Providing such explicit information about interfering transmissions is the topic of an ongoing 3GPP study [RP-130404].

Finally, more radical changes are possible with completely new air interfaces in later releases or standards that can be described as 5G. Besides inter-cell interference cancellation, also intra-cell interference cancellation could be required if the users within a cell are not completely orthogonalized. For example, NTT DoCoMo suggests [SBKN13; BLS+13] to employ a *non-orthogonal multiple access* (NOMA) scheme that schedules two or more users in a non-orthogonal way on identical resources. Instead of (or in addition to) spatial processing, they propose a power loading that exploits pathloss differences between co-scheduled users. To do so, a high-pathloss cell-edge user could be scheduled with a larger fraction of the total power in parallel to a low-pathloss cell-center user who only is assigned a low transmission power. That way, the cell-edge user can still receive its transmission because the intra-cell interference is comparably low, while the cell-center user can first detect and then successively cancel the stronger co-scheduled signal in order to receive the weaker signal that is still good enough in its low-pathloss conditions. Extending the use of cancellation techniques in the downlink puts higher processing demands on the MS, but following *Moore's law* [Mol06], more than enough processing power should be available by the time such new standards are deployed in the field.

6.4. Summary and Contributions

In this final chapter of the thesis, we considered advanced RRM schemes with a special focus on multi-cell aspects. First, we discussed static interference management techniques that aim to reduce interference by restricting the frequency reuse in neighboring cells. Since we could show that, if evaluated properly, such RRM schemes do not offer spectral efficiency gains in most situations, we considered RRM schemes employing a full reuse and advanced spatial processing in the second part of the chapter.

In the first part, we reviewed conventional static reuse schemes including fractional frequency reuse, which combines a full reuse with a reuse 3 at the cell edge. A first analysis of the interference and reuse trade-off in a simplified SISO setup suggested that—in line with a large part of the literature—a reuse 3 could indeed achieve a spectral efficiency gain for some users. But as soon as only basic LTE RRM schemes like PF scheduling or spatial processing are considered, any potential gain disappears. We confirmed this assessment by system-level simulations showing that the PF scheduler even allows for a better trade-off between cell-average and cell-edge rates than FFR.

More promising than static schemes are schemes that dynamically coordinate the system-wide reuse. As our second contribution, we thus proposed a reuse coordination scheme relying on autonomously-deciding base stations. Our method does not need any signaling except for the created interference itself. By applying game theoretic concepts, we showed that it nevertheless converges to a stable system-wide solution. An evaluation of the average interference reduction shows that it cannot reach the optimum interference level, but significantly improves over worst-case and random allocations. Despite its low complexity and stability, the general disadvantages of reuse restrictions likely limit its use to special settings like the examined femtocell case.

Motivated by the flashlight effect from Chapter 5, our next contribution proposed a *neighbor-considerate* modification to TM6 beamforming enforcing a fixed PMI-to-subband mapping in each cell. We showed that despite a loss in diversity, the gains in link adaptation due to the stable interference significantly improve the spectral efficiency.

A more stable interference situation is also an important aspect in our final contribution. We evaluated the performance of different multi-user MIMO schemes with a special emphasis on system-level implications like scheduling and link adaptation. MU-MIMO allows exploiting multiple BS Tx antennas for spatial multiplexing of different users also if individual BS-MS links do not support per-link spatial multiplexing due to channel correlations or too few MS Rx antennas. The considered zero-forcing and PU²RC schemes rely on a channel vector quantization method for CSI feedback whose signaling overhead is identical to that of TM4 and TM6. Our evaluation showed that both ZF and PU²RC can significantly increase the spectral efficiency compared to single-user MIMO with TM4 and TM6. However, ZF requires a more complex user grouping process, suffers from less-precise interference covariance estimation, and needs a higher PHY overhead than PU²RC, which is thus our preferred choice.

We concluded the chapter and the thesis with an outlook to potential future enhancements to LTE-Advanced and emerging trends for 5G networks.

7. Conclusions

This thesis has shown that the downlink spectral efficiency of the current LTE standard, as well as of any potential improvements, depends on many factors and their interplay. Their evaluation requires a careful system-level modeling taking the employed radio resource management schemes, as well as the characteristics of the deployment scenario, the users, and the radio propagation conditions, into account. To this end, we selected the necessary models and provided an open-source simulation tool that, after extensive calibration with reference results from industry-leading 3GPP members, allows an efficient and quantitative performance evaluation. The main contribution of the thesis thus lies in the consistent and holistic performance evaluation of multiple RRM schemes in a comparable way.

The thesis focused on three areas. First, we identified the appropriate scenarios, radio channel, and link-to-system modeling infrastructure and provided an efficient simulation tool named `IMTAphy`. Second, we investigated the performance of fundamental LTE radio resource management schemes taking dependencies between schemes and scenario properties into account. Third, we investigated promising RRM schemes that could increase the spectral efficiency of LTE or LTE-Advanced.

Our conclusions and contributions in the above areas can be summarized as follows:

Downlink System Spectral Efficiency of Fundamental LTE RRM Schemes

We provided a detailed and holistic system-level performance evaluation of the LTE downlink that helps to better understand the spectral efficiency dependencies and trade-offs between RRM schemes and scenarios. We motivated that no single system or cell-edge spectral efficiency figure can be given because the typical deployment scenario and RRM configuration does not exist. Instead, we discussed the performance of different transmission modes and RRM schemes in four distinct IMT-Advanced evaluation scenarios. Thanks to our fully-calibrated simulation tool, all quantitative results, ranging from, for example, a cell spectral efficiency of about 0.8 bit/s/Hz for the simplest SISO round-robin scheduler in the urban micro scenario, to above 10 bit/s/Hz for 4×4 closed-loop MIMO transmission with single-layer interference in the indoor hotspot scenario, are both comparable and fully-reproducible.

Regarding the fundamental LTE RRM schemes, we analyzed in particular the dependency of channel-adaptive scheduling gains on the amount of frequency- and time-domain CQI feedback. Besides outdated channel state feedback and too aggressive scheduling, we showed how fluctuations in the power and spatial structure of the inter-cell interference deteriorate the link adaptation performance. We presented a

solution in the form of outer-loop link adaptation, which is an unstandardized yet essential LTE RRM feature because its use is almost always advantageous.

Another focus was on the evaluation of Tx- and Rx-side spatial processing enabled by multiple antennas. Especially if multiple Rx antennas at the MS are available, large system-level gains can be obtained. However, except for the indoor hotspot scenario, these gains do not stem from higher order ranks of single-user spatial multiplexing. They are rather due to the interference suppression enabled by multiple Rx antennas. Adding more Tx antennas at the BS can also be beneficial. However, since the Tx processing in an FDD system has to be based on limited feedback provided by the MS, and because Tx processing creates higher interference uncertainty, the associated system-level gains are smaller than those from additional Rx antennas. Further, we compared open-loop versus closed-loop MIMO spatial multiplexing as provided by LTE transmission modes TM3 and TM4, respectively. We could show that the benefit of single-user spatial multiplexing, as well as the distribution of the chosen transmission ranks, depend on the deployment scenarios and especially on the use of cross- or co-polarized antenna arrays. Another dependency we have shown to greatly affect the performance of a BS–MS link is the kind of RRM schemes used by interfering base stations. On the one hand, this means that researchers should be aware that when comparing RRM schemes, which are evaluated against their own interference backdrop, a significant part of the observed difference between schemes can stem from the different interference situation and not from the performance on the considered link. On the other hand, this means a base station should also take the system-wide repercussions into account when choosing the transmit strategy for a certain user.

The above-mentioned dependencies motivate our system-level simulation approach and the careful choice of models allowing to obtain meaningful results. Note that, as many RRM mechanisms are not standardized, their performance depends on the individual algorithms implemented in a system. We selected representative implementations, for example, for our schedulers. But it is in the nature of a simulation-based performance evaluation that it is mostly not possible to establish upper bounds on the performance because better algorithms might exist. However, the chosen algorithms can serve as a constructive lower bound on the system performance because their implementation shows how to achieve it.

Selected Advanced RRM Schemes

For our analysis of advanced RRM methods that go beyond the fundamental LTE RRM methods, we picked two directions. First, we considered interference management schemes based on frequency reuse coordination. We could show that a popular static coordination technique called fractional frequency reuse cannot increase the spectral efficiency if standard features like channel-adaptive scheduling and at least two Rx antennas are taken as a baseline. Instead, a full frequency reuse should be used in almost all scenarios. One exception might be the operation of so-called femtocells. For such scenarios, we showed that a dynamic interference coordination method can converge to a stable allocation even if decisions are made autonomously in each cell. By

comparing to an optimal allocation as a benchmark, we evaluated the interference reduction achieved by the proposed scheme working without explicit signaling between the cells, which might be a requirement in such unplanned deployments.

As the second direction, we investigated predictable beamforming and multi-user MIMO transmission schemes that address two of the challenges we identified with respect to the fundamental RRM schemes. Predictable beamforming allows to improve the beamforming performance by making the interference less fluctuating, which helps the link adaptation. MU-MIMO primarily aims at increasing the spatial multiplexing gains by distributing the transmitted layers over multiple users, which overcomes the low channel ranks we observed for single-user spatial multiplexing. Especially the PU²RC MU-MIMO variant, in addition, also makes the intra- and inter-cell interference more predictable. This helps PU²RC to exceed the performance of the zero-forcing reference scheme that we investigated as well. Overall, MU-MIMO is clearly beneficial in the investigated urban micro scenario, where the performance of single-user MIMO schemes is limited due to the highly correlated BS–MS links.

Since the considered advanced RRM schemes are meant as extensions to the basic LTE operation requiring no or only minimal changes to the standard, they do not replace the fundamental RRM schemes we discussed, but rather extend them. Building upon our analysis of fundamental RRM schemes, and having a detailed channel model as the foundation, thus helped us to avoid meaningless conclusions. For example, publications in the literature that analyze fractional frequency reuse strategies without the mentioned fundamental RRM schemes might suggest FFR benefits that would not be realized in practice. Also, evaluations of MU-MIMO strategies that do not consider limited feedback, the necessary overhead for demodulation reference signals, and the impacts on the interference estimation at the receiver and on the link adaptation performance, might be insightful. However, they are not sufficient to select candidate schemes for future standards.

Channel and Link-to-System Models for System-Level Simulation

Motivated by the observed high dependency of the link and system performance on the scenario and radio channel conditions as well as on the receiver performance, we selected suitable models according to three goals: First, the models need to reflect real-world deployment scenarios, radio propagation characteristics, and receivers well-enough to allow deriving meaningful quantitative results. Second, the modeling assumptions and model parameters need to be chosen as to allow comparison with results from third parties. Third, the model needs an abstraction level that allows an efficient implementation in order to run the required simulations with a feasible amount of compute resources.

The models presented in this thesis and implemented in `IMTAphy` fulfill all of the above goals. The scenario modeling and the geometry-based spatial channel model follow the ITU guidelines for IMT-Advanced, which have become a de-facto standard. The results are thus highly comparable. The channel models can also be expected to

represent reality well-enough because they are based on comprehensive measurement campaigns. The link-to-system model also follows standard modeling assumptions from the literature. All model implementations in IMT_{Aphy} are extensively calibrated as outlined in Chapters 3 and 4. In addition, the complete source code of the model implementations, as well as of all RRM schemes evaluated in this thesis, is made available as open-source on the Internet. This allows interested third parties to validate the implementation, reproduce the results, and extend it. Finally, the presented analysis of the channel model led to a very efficient implementation in IMT_{Aphy}, which allowed us to conduct more than 200 000 simulation runs for this thesis.

Outlook

The development of cellular systems does not stop with LTE-Advanced as 5G is already at the horizon. In Section 6.3 we outlined a number of promising future enhancements, which we do not want to repeat here. However, we want to sketch a few implications of some of the main trends (e.g., mm-wave, massive and network MIMO, device-to-device, and offloading to WLAN) for future system-level modeling needs.

Going forward, it is perceivable that system-level modeling approaches will have to be extended both in breadth and in depth. They have to be extended in breadth to accommodate the foreseeable tighter integration of the traditional BS–MS cellular access link with other radio access technologies that aim at offloading, for example, by using WLAN access points or direct D2D links between mobiles. They also have to be extended in depth. The current link-to-system model is tailored to linear receivers, but to achieve higher link efficiencies, more advanced receiver types employing, for example, maximum likelihood detection or interference cancellation will be used in the future. Unless suitable link abstraction models can be found, one foreseeable alternative is to perform traditional link-level simulations on a system level taking the full system-wide interference situation and RRM interdependencies into account. The highly parallel nature of per-link evaluations, as well as the ever-increasing computational resources provided by *Moore's law*, will soon enable such a simulation approach, which can overcome the limitations of current link-to-system models.

Also, the channel models have to be extended in order to cope with higher and larger frequency bandwidths in the mm-wave region and to accommodate the currently unsupported link between two mobiles, which is relevant for D2D and asynchronous TDD deployments. In addition to supporting new link types, also the correlations and synchronization properties of the links from one user to multiple transmission points should be modeled in more detail to allow a better evaluation of CoMP transmission techniques. Furthermore, more focus should be put on RF aspects, especially on the mobile terminal side. For example, the current antenna modeling assumptions are over-simplified and do not reflect the challenges of placing sufficiently-uncorrelated antennas on a small mobile device.

In any case, new modeling guidelines and calibration approaches will be required to cope with the complexity and to enable meaningful and comparable simulation results allowing future standardization bodies to make well-founded technology decisions.

A. Additional Simulation Results for Chapter 5

In this section, we provide a few additional simulation results for the evaluation of fundamental RRM schemes in Chapter 5. The tables and figures printed here are referenced in the main text of the thesis.

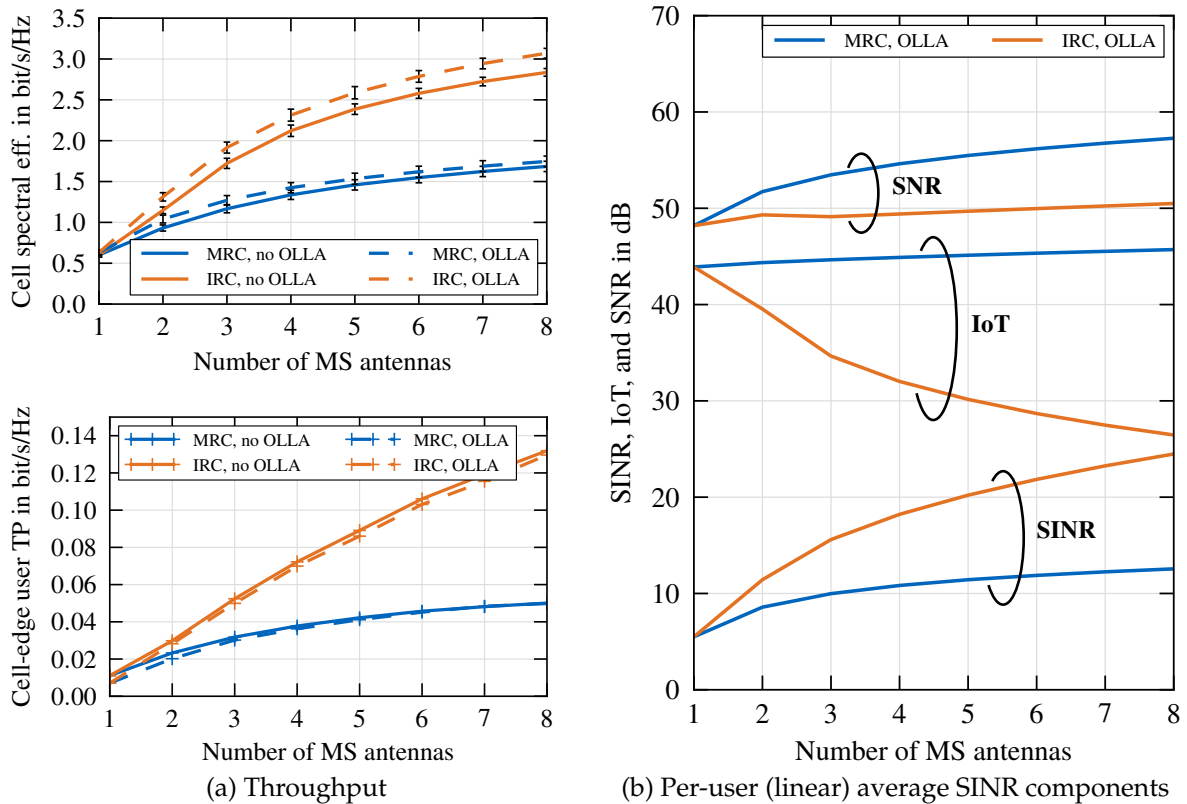


Figure A.1.: Throughput (TP) gains in terms of system spectral efficiency (top left) and cell-edge user throughput (bottom left) depending on the number of Rx antennas and the receive filter and link adaptation scheme in the UMa scenario with the underlying post-combining mean per-user (linear) average SINRs separated into mean SNR and mean IoT

A. Additional Simulation Results for Chapter 5

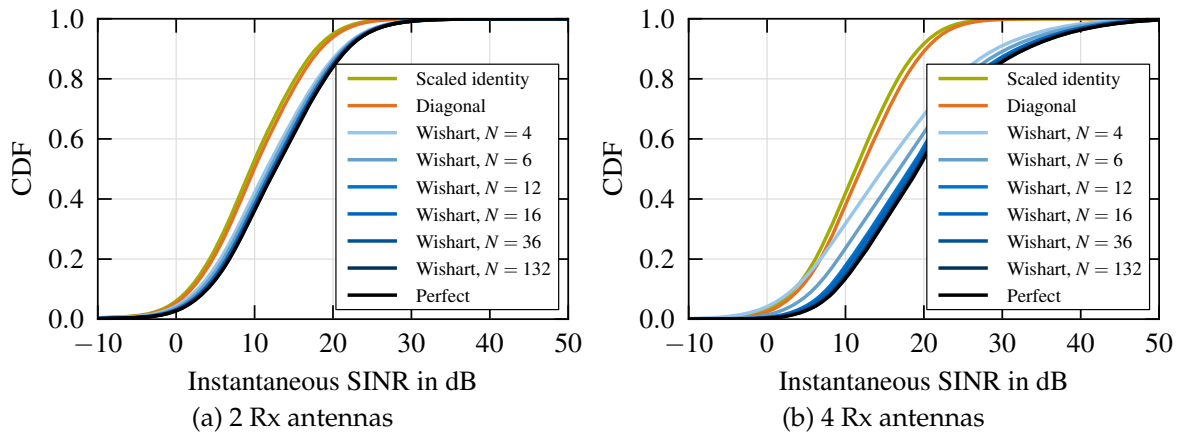


Figure A.2.: Instantaneous SINR achieved by the 4 Rx MMSE(-IRC) receiver using different IPN covariance estimation methods in the UMi scenario

Table A.1.: Frequency of PMI selection, 2 Tx antennas

PMI	RMa		UMa		UMi	
	A	C	A	C	A	C
0	25.8%	28.3%	25.4%	34.5%	25.4%	32.8%
1	24.5%	12.5%	25.0%	8.2%	24.3%	9.9%
2	25.3%	29.7%	24.8%	29.5%	25.0%	28.5%
3	24.4%	29.4%	24.9%	27.9%	25.3%	28.9%

Table A.2.: Frequency of PMI selection, 4 Tx antennas

PMI	RMa		UMa		UMi		PMI	RMa		UMa		UMi	
	A	C	A	C	A	C		A	C	A	C	A	C
0	11.8%	13.6%	10.5%	16.3%	9.2%	15.3%	8	1.2%	0.6%	2.5%	1.3%	4.0%	1.4%
1	11.3%	14.9%	9.9%	14.5%	9.0%	14.1%	9	1.3%	0.5%	2.5%	0.5%	3.8%	0.6%
2	10.7%	1.0%	9.7%	1.1%	7.6%	1.2%	10	1.3%	0.4%	2.5%	0.5%	3.8%	0.6%
3	10.9%	14.6%	9.7%	14.1%	8.1%	14.3%	11	1.2%	0.4%	2.5%	0.4%	4.0%	0.6%
4	11.2%	14.2%	9.3%	16.2%	7.8%	15.1%	12	1.4%	0.5%	2.8%	0.6%	4.4%	1.1%
5	11.2%	12.1%	9.6%	8.9%	7.9%	9.5%	13	1.5%	0.2%	3.0%	0.2%	4.7%	0.3%
6	11.0%	12.1%	9.6%	8.3%	8.4%	9.4%	14	1.4%	0.2%	3.0%	0.2%	4.6%	0.3%
7	11.3%	14.3%	10.1%	16.2%	8.3%	15.1%	15	1.3%	0.5%	2.9%	0.7%	4.5%	1.1%

Table A.3.: TM6 system spectral efficiencies in bit/s/Hz depending on the receiver type, antenna configuration, and antenna spacing in the UMa scenario

Receiver / array type	MIMO link ($T_x \times R_x$) antenna configuration											
	1×1	1×2	1×3	1×4	2×1	2×2	2×3	2×4	4×1	4×2	4×3	4×4
MRC, config C	0.63	1.04	1.27	1.42	0.85	1.27	1.48	1.63	1.11	1.51	1.72	1.86
MRC, config A	0.63	1.04	1.27	1.42	0.75	1.14	1.35	1.49	0.91	1.30	1.50	1.62
IRC, config C	0.63	1.31	1.92	2.31	0.85	1.55	2.12	2.48	1.11	1.82	2.34	2.66
IRC, config A	0.63	1.31	1.92	2.31	0.75	1.42	1.99	2.36	0.91	1.58	2.12	2.46




Figure A.3.: Cell spectral efficiency and normalized cell-edge user throughput achieved by MMSE linear receiver with different IPN covariance matrix estimation models in the UMi scenario. See Table 5.9 for simulation parameters.

Table A.4.: TM6 cell-edge user throughput in bit/s/Hz depending on the receiver type, antenna configuration, and antenna spacing in the UMa scenario

Receiver / array type	MIMO link (Tx × Rx) antenna configuration											
	1 × 1	1 × 2	1 × 3	1 × 4	2 × 1	2 × 2	2 × 3	2 × 4	4 × 1	4 × 2	4 × 3	4 × 4
MRC, config C	0.007	0.020	0.030	0.036	0.010	0.027	0.037	0.045	0.016	0.037	0.049	0.057
MRC, config A	0.007	0.020	0.030	0.036	0.008	0.025	0.034	0.041	0.010	0.030	0.042	0.049
IRC, config C	0.007	0.028	0.050	0.070	0.010	0.037	0.062	0.082	0.016	0.050	0.078	0.098
IRC, config A	0.007	0.028	0.050	0.070	0.008	0.034	0.056	0.074	0.010	0.039	0.063	0.081

Table A.5.: System cell-edge user spectral efficiency of TM3 and TM4 compared to TM1 depending on scenario and antenna array configurations

Scenario and array config	Avg. spectral efficiency in bit/s/Hz						Avg. cell-edge user throughput in bit/s/Hz						
	MMSE-MRC			MMSE-IRC			MMSE-MRC			MMSE-IRC			
	TM1	TM3	TM4	TM1	TM3	TM4	TM1	TM3	TM4	TM1	TM3	TM4	
RMa	C↔C	1.61	1.54	2.14	2.36	2.15	2.78	0.045	0.034	0.070	0.084	0.066	0.107
	A↔C	n/a	1.60	2.06	n/a	2.16	2.63	n/a	0.031	0.060	n/a	0.059	0.091
	E↔E	n/a	1.93	2.49	n/a	2.42	3.03	n/a	0.039	0.059	n/a	0.057	0.082
	D↔E	n/a	2.04	2.44	n/a	2.42	2.89	n/a	0.040	0.054	n/a	0.056	0.073
InH	C↔C	3.18	5.68	5.57	3.67	6.63	6.59	0.141	0.081	0.108	0.314	0.259	0.239
	A↔C	n/a	6.16	5.99	n/a	6.86	6.77	n/a	0.089	0.110	n/a	0.233	0.223
	E↔E	n/a	7.20	7.12	n/a	7.56	7.35	n/a	0.103	0.109	n/a	0.255	0.271
	D↔E	n/a	7.33	7.28	n/a	7.62	7.45	n/a	0.104	0.118	n/a	0.255	0.256

A. Additional Simulation Results for Chapter 5




Figure A.4.: Geographic distribution of PMIs in the UMa scenario (low LoS prob.)

Table A.6.: Overview of system spectral efficiency and cell-edge user throughput for open- and closed-loop MIMO compared to TM1 depending on scenario and antenna array configurations

Scenario and array config	Avg. spectral efficiency in bit/s/Hz						Avg. cell-edge user throughput in bit/s/Hz						
	MMSE-MRC			MMSE-IRC			MMSE-MRC			MMSE-IRC			
	TM1	TM3	TM4	TM1	TM3	TM4	TM1	TM3	TM4	TM1	TM3	TM4	
UMa	C \leftrightarrow C	1.52	1.44	2.05	2.32	2.02	2.66	0.041	0.029	0.058	0.071	0.052	0.084
	A \leftrightarrow C	n/a	1.50	1.82	n/a	1.95	2.27	n/a	0.029	0.049	n/a	0.046	0.068
	E \leftrightarrow E	n/a	1.66	2.09	n/a	2.08	2.55	n/a	0.033	0.047	n/a	0.045	0.063
	D \leftrightarrow E	n/a	1.68	1.94	n/a	2.01	2.34	n/a	0.032	0.042	n/a	0.042	0.055
UMi	C \leftrightarrow C	1.85	1.59	2.20	2.76	2.34	2.91	0.050	0.027	0.062	0.090	0.053	0.092
	A \leftrightarrow C	n/a	1.71	2.11	n/a	2.26	2.65	n/a	0.029	0.054	n/a	0.050	0.078
	E \leftrightarrow E	n/a	2.05	2.54	n/a	2.57	3.10	n/a	0.038	0.063	n/a	0.057	0.084
	D \leftrightarrow E	n/a	2.12	2.48	n/a	2.50	2.93	n/a	0.039	0.058	n/a	0.055	0.075

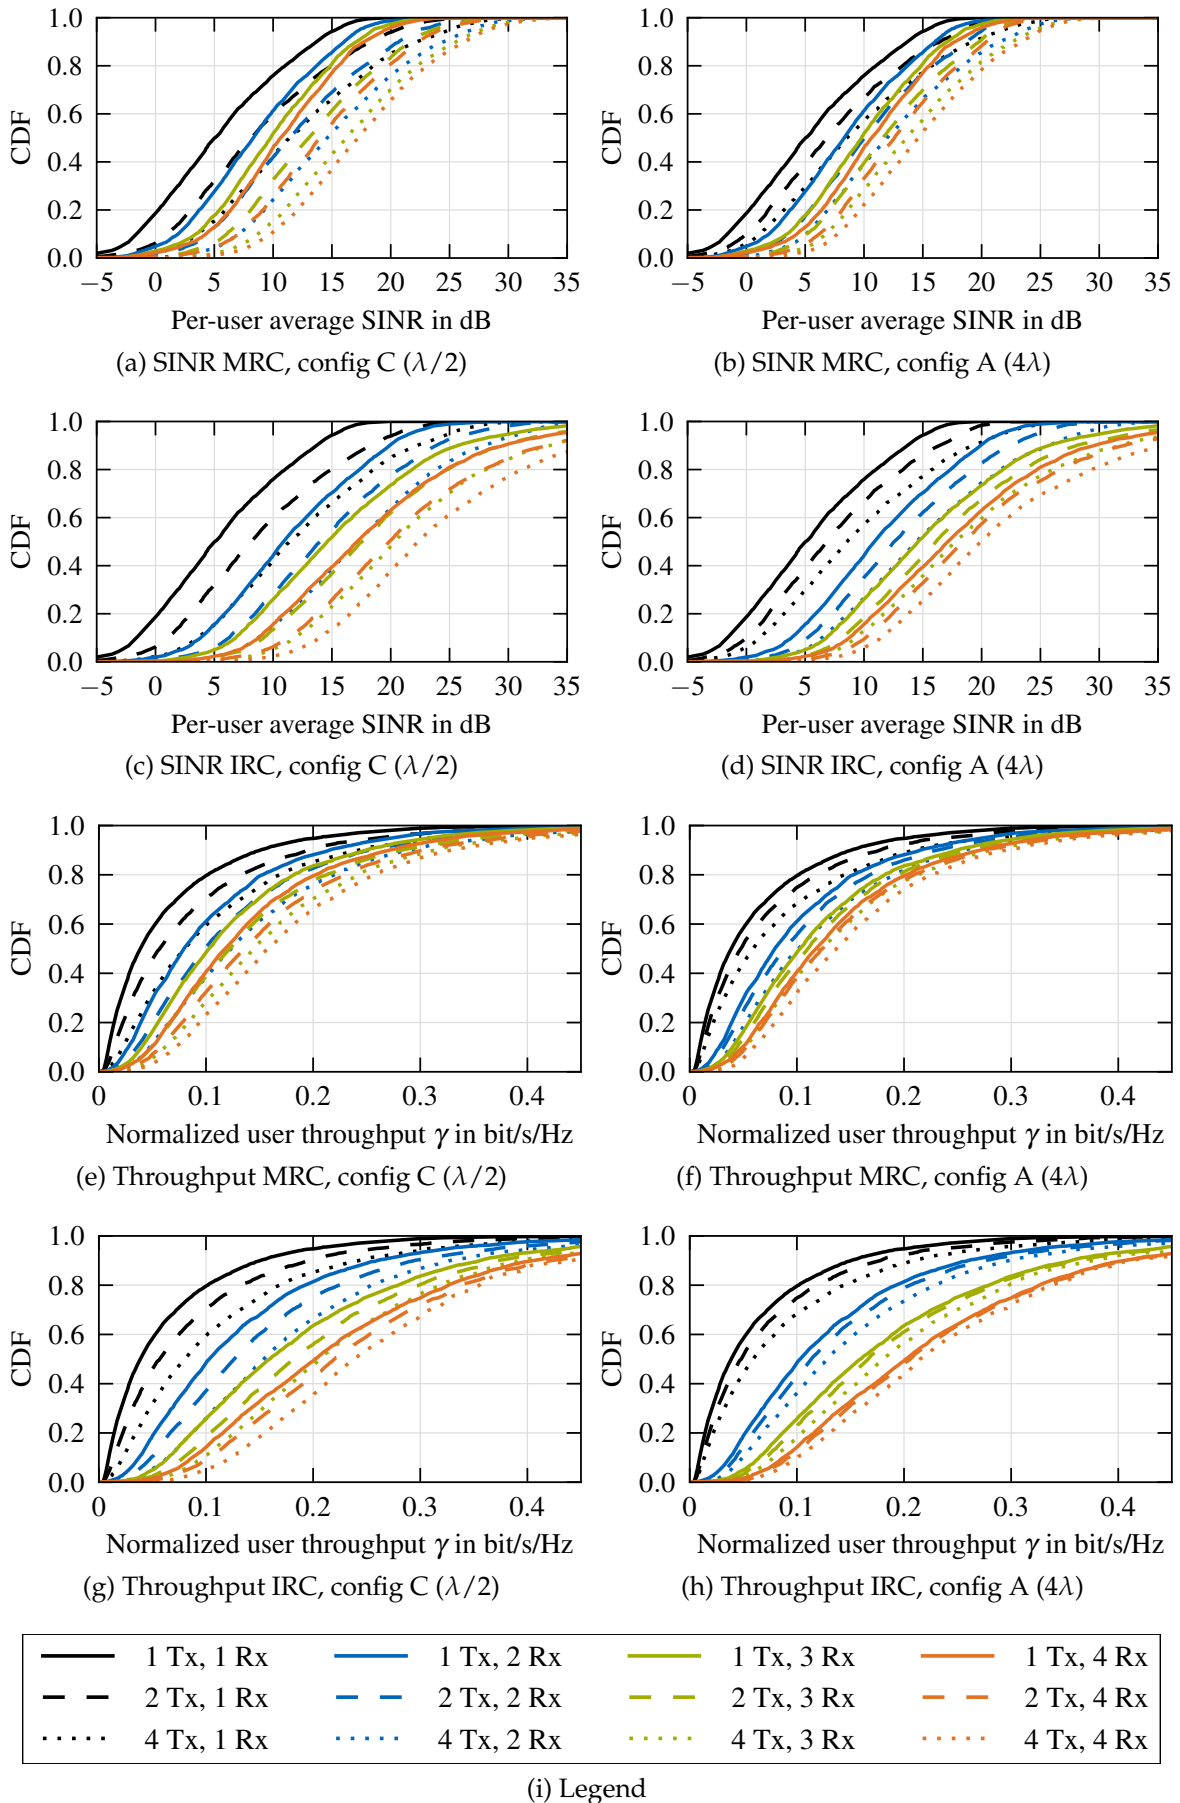


Figure A.5.: Performance of TM6 in the UMa scenario

A. Additional Simulation Results for Chapter 5



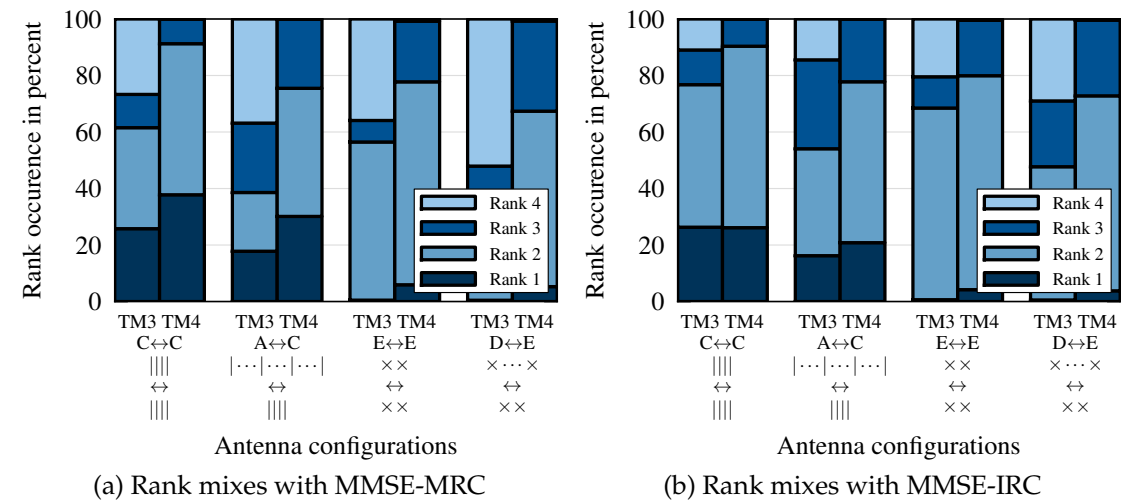


Figure A.6.: Impact of 4×4 MIMO antenna configurations on the rank mixes and per-user throughput distribution in the UMa scenario; four antenna array combinations (BS \leftrightarrow MS), each for TM3 and TM4 as well as TM1 as a reference for the throughput

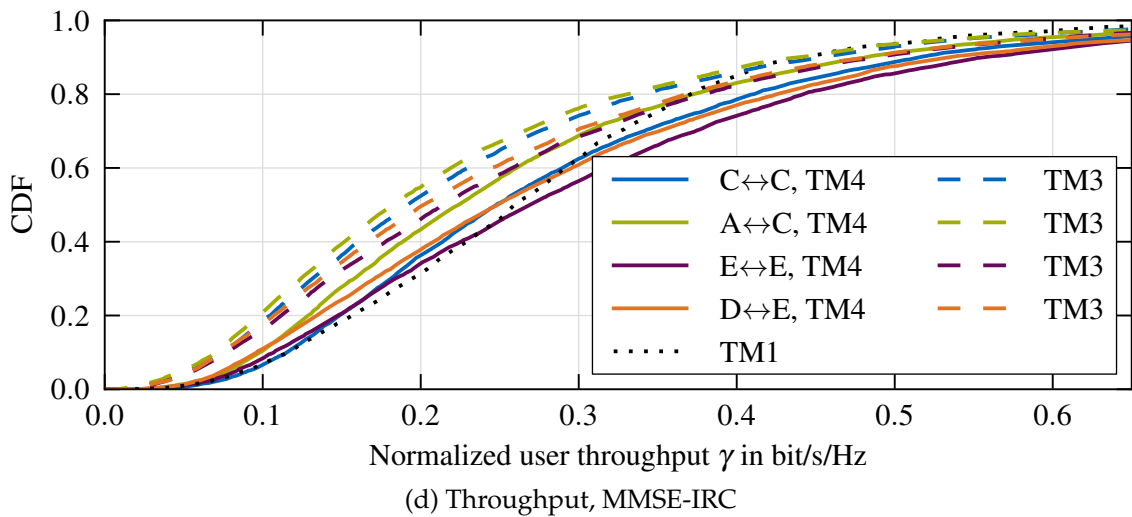

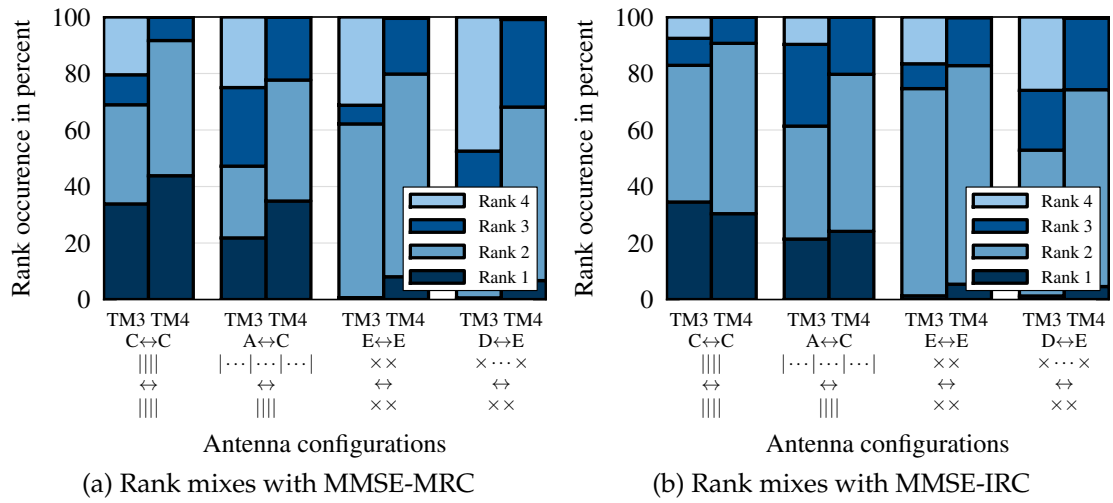


Figure A.7.: Impact of 4×4 MIMO antenna configurations on the rank mixes and per-user throughput distribution in the UMi scenario; four antenna array combinations (BS \leftrightarrow MS), each for TM3 and TM4 and TM1 as a reference for the throughput

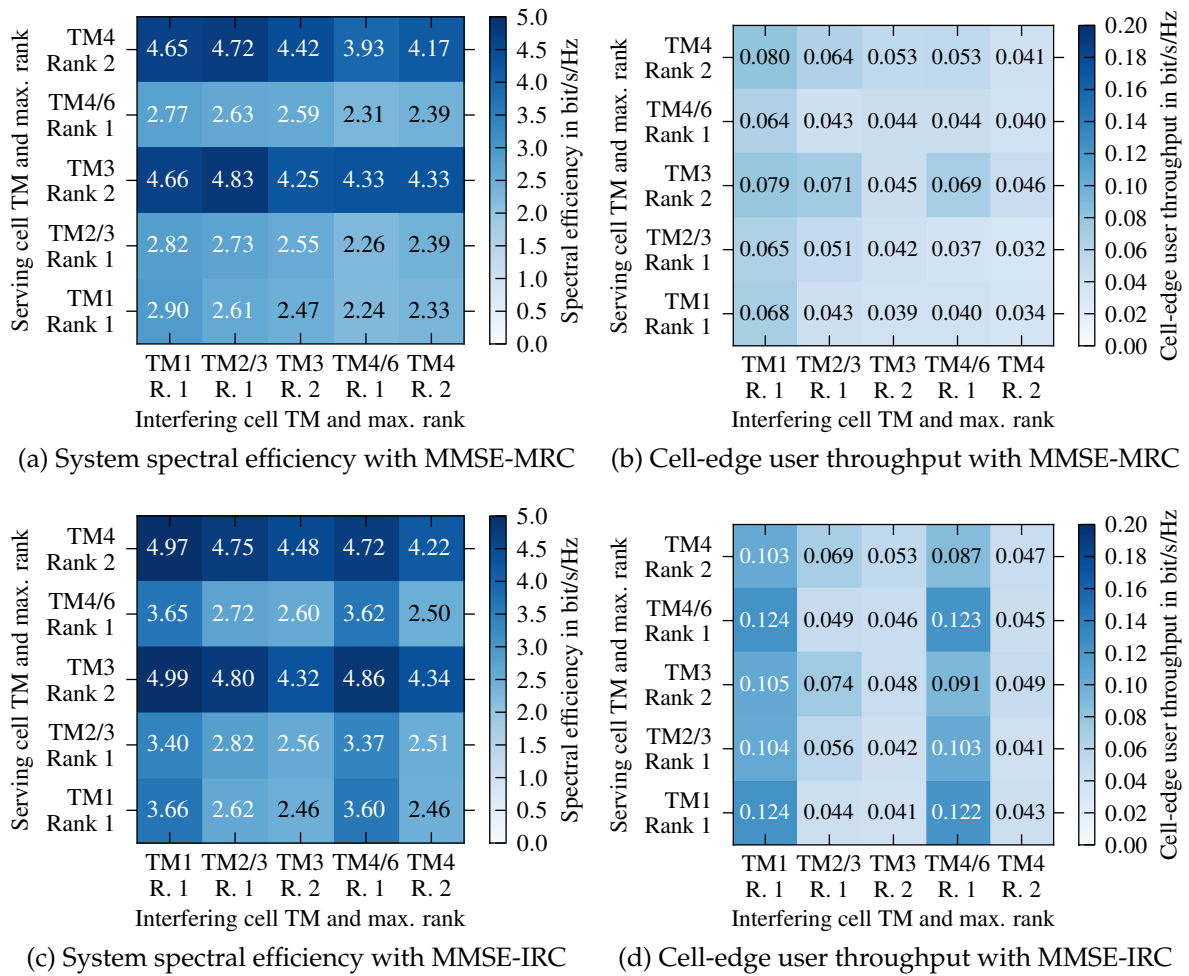


Figure A.8.: System and cell-edge user spectral efficiency in the InH scenario with a 4×4 antenna configuration (E-E), evaluating those 10% of users with TM and max. rank according to the y-axis while the other 90% of users cause interference with TMs and max. ranks as indicated on the x-axis

B. Codebooks from the LTE Standard

CQI Feedback

Table B.1.: LTE CQI definition [36.213, Sec. 7.2.3]

CQI index	Modulation scheme	Code rate	Net rate $r(\text{CQI})$ in bit/s/Hz
0	out of range		
1	QPSK	0.076	0.1523
2	QPSK	0.117	0.2344
3	QPSK	0.188	0.3770
4	QPSK	0.301	0.6016
5	QPSK	0.438	0.8770
6	QPSK	0.588	1.1758
7	16-QAM	0.369	1.4766
8	16-QAM	0.479	1.9141
9	16-QAM	0.602	2.4063
10	64-QAM	0.455	2.7305
11	64-QAM	0.554	3.3223
12	64-QAM	0.650	3.9023
13	64-QAM	0.754	4.5234
14	64-QAM	0.853	5.1152
15	64-QAM	0.926	5.5547

Precoding Matrix Definitions for Open-Loop Operation

Table B.2.: LTE open-loop spatial multiplexing precoding factors

Layers M	\mathbf{U}	$\mathbf{D}(i)$
2	$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & e^{-j2\pi/2} \end{pmatrix}$	$\begin{pmatrix} 1 & 0 \\ 0 & e^{-j2i\pi/2} \end{pmatrix}$
3	$\frac{1}{\sqrt{3}} \begin{pmatrix} 1 & 1 & 1 \\ 1 & e^{-j2\pi/3} & e^{-j4\pi/3} \\ 1 & e^{-j4\pi/3} & e^{-j8\pi/3} \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{-j2i\pi/3} & 0 \\ 0 & 0 & e^{-j4i\pi/3} \end{pmatrix}$
4	$\frac{1}{\sqrt{4}} \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & e^{-j2\pi/4} & e^{-j4\pi/4} & e^{-j6\pi/4} \\ 1 & e^{-j4\pi/4} & e^{-j8\pi/4} & e^{-j12\pi/4} \\ 1 & e^{-j6\pi/4} & e^{-j12\pi/4} & e^{-j18\pi/4} \end{pmatrix}$	$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & e^{-j2i\pi/4} & 0 & 0 \\ 0 & 0 & e^{-j4i\pi/4} & 0 \\ 0 & 0 & 0 & e^{-j6i\pi/4} \end{pmatrix}$

Precoding Matrix Definitions for Closed-Loop Operation

Codebooks as defined in [36.211, pp. 6.3.4.2.3-1] for 2 Tx antennas and in [36.211, pp. 6.3.4.2.3-2] for 4 Tx antennas. In the last column of the tables, the linear array gains (radiation patterns) realized by a 2 or 4 antenna linear antenna array along the y-axis with an antenna spacing of $\lambda/2$ are plotted. With this layout, the array broadside is pointing towards the right into the 0° direction. Note that for a base station in a sectorized deployment the sector pattern would be superimposed covering an azimuth range from approximately -60 to $+60$ degrees. The first layer is plotted in blue, the second layer in red, the third layer in green, and the fourth layer in cyan. The sum of all layers is plotted in black.

The 4 Tx antenna codebook offers precoding matrices of size $N_{\text{Tx}} \times \text{rank}$ which are taken from a master matrix \mathbf{P} [1234]. In the table the chosen columns per rank are indicated by the numbers in brackets. Note that the plot in the last column shows the radiation patterns of the rank 4 precoding matrix, which can have a different ordering of columns as the master matrix \mathbf{P} [1234]. All matrices are defined in a way by the standard that assures that the radiation power is equally distributed over layers and transmit antennas.

Table B.3.: LTE 4 Tx Codebook, PMIs 0–3 (column 1: blue, column 2: red, column 3: green, column 4: cyan, total: black)

PMI	$\mathbf{P}[1234]$	Rank 1	Rank 2	Rank 3	Rank 4	Plot of Rank 4 Matrix
0	$\frac{1}{2} \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & 1 & -1 \\ 1 & -1 & -1 & 1 \end{pmatrix}$	$\mathbf{P}[1]$	$\frac{1}{\sqrt{2}} \mathbf{P}[14]$	$\frac{1}{\sqrt{3}} \mathbf{P}[124]$	$\frac{1}{2} \mathbf{P}[1234]$	
1	$\frac{1}{2} \begin{pmatrix} 1 & -j & -1 & j \\ j & 1 & j & 1 \\ -1 & -j & 1 & j \\ -j & 1 & -j & 1 \end{pmatrix}$	$\mathbf{P}[1]$	$\frac{1}{\sqrt{2}} \mathbf{P}[12]$	$\frac{1}{\sqrt{3}} \mathbf{P}[123]$	$\frac{1}{2} \mathbf{P}[1234]$	
2	$\frac{1}{2} \begin{pmatrix} 1 & -1 & 1 & -1 \\ -1 & 1 & 1 & -1 \\ 1 & 1 & 1 & 1 \\ -1 & -1 & 1 & 1 \end{pmatrix}$	$\mathbf{P}[1]$	$\frac{1}{\sqrt{2}} \mathbf{P}[12]$	$\frac{1}{\sqrt{3}} \mathbf{P}[123]$	$\frac{1}{2} \mathbf{P}[3214]$	
3	$\frac{1}{2} \begin{pmatrix} 1 & j & -1 & -j \\ -j & 1 & -j & 1 \\ -1 & j & 1 & -j \\ j & 1 & 1 & j \end{pmatrix}$	$\mathbf{P}[1]$	$\frac{1}{\sqrt{2}} \mathbf{P}[12]$	$\frac{1}{\sqrt{3}} \mathbf{P}[123]$	$\frac{1}{2} \mathbf{P}[3214]$	

B. Codebooks from the LTE Standard

Table B.4.: LTE 4 Tx Codebook, PMIs 4–7 (column 1: blue, column 2: red, column 3: green, column 4: cyan, total: black)

PMI	$P_{[1234]}$	Rank 1	Rank 2	Rank 3	Rank 4	Plot of Rank 4 Matrix
4	$\frac{1}{2} \begin{pmatrix} 1 & \frac{1-j}{\sqrt{2}} & -j & \frac{-1-j}{\sqrt{2}} \\ \frac{1+j}{\sqrt{2}} & 1 & \frac{-1+j}{\sqrt{2}} & j \\ j & \frac{1+j}{\sqrt{2}} & 1 & j \\ \frac{-1+j}{\sqrt{2}} & -j & \frac{-1-j}{\sqrt{2}} & 1 \end{pmatrix}$	$P_{[1]}$	$\frac{1}{\sqrt{2}}P_{[14]}$	$\frac{1}{\sqrt{3}}P_{[124]}$	$\frac{1}{2}P_{[1234]}$	
5	$\frac{1}{2} \begin{pmatrix} 1 & \frac{-1-j}{\sqrt{2}} & j & \frac{1-j}{\sqrt{2}} \\ \frac{-1+j}{\sqrt{2}} & 1 & \frac{1+j}{\sqrt{2}} & -j \\ j & \frac{1+j}{\sqrt{2}} & 1 & \frac{1-j}{\sqrt{2}} \\ \frac{1+j}{\sqrt{2}} & -j & \frac{-1-j}{\sqrt{2}} & 1 \end{pmatrix}$	$P_{[1]}$	$\frac{1}{\sqrt{2}}P_{[14]}$	$\frac{1}{\sqrt{3}}P_{[124]}$	$\frac{1}{2}P_{[1234]}$	
6	$\frac{1}{2} \begin{pmatrix} 1 & \frac{-1-j}{\sqrt{2}} & -j & \frac{1+j}{\sqrt{2}} \\ \frac{-1+j}{\sqrt{2}} & 1 & \frac{1+j}{\sqrt{2}} & -j \\ -j & \frac{1+j}{\sqrt{2}} & 1 & \frac{1+j}{\sqrt{2}} \\ \frac{1+j}{\sqrt{2}} & j & \frac{1-j}{\sqrt{2}} & 1 \end{pmatrix}$	$P_{[1]}$	$\frac{1}{\sqrt{2}}P_{[13]}$	$\frac{1}{\sqrt{3}}P_{[134]}$	$\frac{1}{2}P_{[1324]}$	
7	$\frac{1}{2} \begin{pmatrix} 1 & \frac{1-j}{\sqrt{2}} & j & \frac{-1+j}{\sqrt{2}} \\ \frac{1-j}{\sqrt{2}} & 1 & \frac{-1+j}{\sqrt{2}} & -j \\ -j & \frac{-1+j}{\sqrt{2}} & 1 & \frac{-1+j}{\sqrt{2}} \\ \frac{-1-j}{\sqrt{2}} & -j & \frac{-1-j}{\sqrt{2}} & 1 \end{pmatrix}$	$P_{[1]}$	$\frac{1}{\sqrt{2}}P_{[13]}$	$\frac{1}{\sqrt{3}}P_{[134]}$	$\frac{1}{2}P_{[1324]}$	

Table B.5.: LTE 4 Tx Codebook, PMIs 8–11 (column 1: blue, column 2: red, column 3: green, column 4: cyan, total: black)

PMI	$\mathbf{P}[1234]$	Rank 1	Rank 2	Rank 3	Rank 4	Plot of Rank 4 Matrix
8	$\frac{1}{2} \begin{pmatrix} 1 & 1 & 1 & -1 \\ 1 & 1 & 1 & 1 \\ -1 & 1 & 1 & -1 \\ -1 & 1 & 1 & 1 \end{pmatrix}$	$\mathbf{P}[1]$	$\frac{1}{\sqrt{2}}\mathbf{P}[12]$	$\frac{1}{\sqrt{3}}\mathbf{P}[124]$	$\frac{1}{2}\mathbf{P}[1234]$	
9	$\frac{1}{2} \begin{pmatrix} 1 & -j & 1 & -j \\ j & 1 & -j & -1 \\ 1 & j & 1 & j \\ j & -1 & -j & 1 \end{pmatrix}$	$\mathbf{P}[1]$	$\frac{1}{\sqrt{2}}\mathbf{P}[14]$	$\frac{1}{\sqrt{3}}\mathbf{P}[134]$	$\frac{1}{2}\mathbf{P}[1234]$	
10	$\frac{1}{2} \begin{pmatrix} 1 & -1 & -1 & 1 \\ -1 & 1 & -1 & 1 \\ -1 & -1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix}$	$\mathbf{P}[1]$	$\frac{1}{\sqrt{2}}\mathbf{P}[13]$	$\frac{1}{\sqrt{3}}\mathbf{P}[123]$	$\frac{1}{2}\mathbf{P}[1324]$	
11	$\frac{1}{2} \begin{pmatrix} 1 & j & 1 & j \\ -j & 1 & -j & -1 \\ 1 & -j & 1 & -j \\ -j & -1 & -j & 1 \end{pmatrix}$	$\mathbf{P}[1]$	$\frac{1}{\sqrt{2}}\mathbf{P}[13]$	$\frac{1}{\sqrt{3}}\mathbf{P}[134]$	$\frac{1}{2}\mathbf{P}[1324]$	

B. Codebooks from the LTE Standard

Table B.6.: LTE 4 Tx Codebook, PMIs 12–15 (column 1: blue, column 2: red, column 3: green, column 4: cyan, total: black)

PMI	$P_{[1234]}$	Rank 1	Rank 2	Rank 3	Rank 4	Plot of Rank 4 Matrix
12	$\frac{1}{2} \begin{pmatrix} 1 & 1 & 1 & -1 \\ 1 & 1 & -1 & 1 \\ -1 & -1 & 1 & 1 \end{pmatrix}$	$P_{[1]}$	$\frac{1}{\sqrt{2}}P_{[12]}$	$\frac{1}{\sqrt{3}}P_{[123]}$	$\frac{1}{2}P_{[1234]}$	
13	$\frac{1}{2} \begin{pmatrix} 1 & 1 & -1 & 1 \\ 1 & 1 & 1 & -1 \\ -1 & -1 & 1 & 1 \end{pmatrix}$	$P_{[1]}$	$\frac{1}{\sqrt{2}}P_{[13]}$	$\frac{1}{\sqrt{3}}P_{[123]}$	$\frac{1}{2}P_{[1324]}$	
14	$\frac{1}{2} \begin{pmatrix} 1 & -1 & 1 & 1 \\ -1 & 1 & 1 & 1 \\ 1 & 1 & -1 & -1 \end{pmatrix}$	$P_{[1]}$	$\frac{1}{\sqrt{2}}P_{[13]}$	$\frac{1}{\sqrt{3}}P_{[123]}$	$\frac{1}{2}P_{[3214]}$	
15	$\frac{1}{2} \begin{pmatrix} 1 & -1 & -1 & -1 \\ -1 & 1 & -1 & -1 \\ -1 & -1 & 1 & -1 \end{pmatrix}$	$P_{[1]}$	$\frac{1}{\sqrt{2}}P_{[12]}$	$\frac{1}{\sqrt{3}}P_{[123]}$	$\frac{1}{2}P_{[1234]}$	

Table B.7.: LTE 2 Tx Codebook, PMIs 0–3 (column 1: blue, column 2: red, total: black)

PMI	Rank 1	Rank 2	Plot Precoding Matrix
0	$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix}$		
1	$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -1 \end{pmatrix}$	$\frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$	
2	$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ j \end{pmatrix}$	$\frac{1}{2} \begin{pmatrix} 1 & 1 \\ j & -j \end{pmatrix}$	
3	$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -j \end{pmatrix}$		

C. Acronyms

%ile	Percentile
16-QAM	Quadrature Amplitude Modulation with constellation size of 16
64-QAM	Quadrature Amplitude Modulation with constellation size of 64
2D	Two-Dimensional
2G	Second Generation
3D	Three-Dimensional
3G	Third Generation
5G	Fifth Generation
3GPP	3rd Generation Partnership Project
ABS	Almost Blank Subframes
ACK	Acknowledgement
AM	Acknowledged Mode
AMPS	Advanced Mobile Phone Service
AoA	Angle of Arrival
AoD	Angle of Departure
ARQ	Automatic Repeat Request
ASA	Angle Spread of Arrival (AoA spread)
ASD	Angle Spread of Departure (AoD spread)
AWGN	Additive White Gaussian Noise
BICM	Bit-Interleaved Coded Modulation
BLER	Block Error Ratio
BS	Base Station
CATR	China Academy of Telecommunication Research
CC	Chase Combining
CCDF	Complementary Cumulative Distribution Function
CDF	Cumulative Distribution Function
CDMA	Code-Division Multiple Access
CI	Column Indicator
CIR	Channel Impulse Response
CLSM	Closed-Loop Spatial Multiplexing
CMCC	China Mobile Communications Corporation
CoMP	Coordinated Multi-Point
CP	Cyclic Prefix
CPU	Central Processing Unit
CQI	Channel Quality Indicator
CRC	Cyclic Redundance Check
CRS	Cell-Specific Reference Signals
CS/CB	Coordinated Scheduling / Coordinated Beamforming

CSI	Channel State Information
CSIR	Channel State Information at the Receiver
CSIT	Channel State Information at the Transmitter
CSI-RS	CSI-Reference Signals
CTF	Channel Transfer Function
D2D	Device-to-Device
DCI	Downlink Control Information
DFT	Discrete Fourier Transform
DL	Downlink
DM-RS	Demodulation Reference Signals
DPC	Dirty Paper Coding
DPB	Dynamic Point Blanking
DPS	Dynamic Point Selection
DS	Delay Spread
DTX	Discontinuous Transmission
EDGE	Enhanced Data Rates for GSM Evolution
EESM	Exponential Effective SINR Mapping
eICIC	Enhanced Inter-Cell Interference Coordination
eNB	Evolved NodeB
EPC	Evolved Packet Core
EPA	Extended Pedestrian A
ePDCCH	Enhanced Physical Downlink Control Channel
ETU	Extended Typical Urban
E-UTRAN	Evolved Universal Terrestrial Radio Access Network
EVA	Extended Vehicular A
EVM	Error Vector Magnitude
FDD	Frequency-Division Duplexing
FDMA	Frequency-Division Multiple Access
FFR	Fractional Frequency Reuse
FFT	Fast Fourier Transform
FIR	Finite Impulse Response
GPRS	General Packet Radio Service
GSM	Global System for Mobile Communication
HARQ	Hybrid Automatic Repeat Request
HII	High Interference Indicator
HSPA	High Speed Packet Access
ICIC	Inter-Cell Interference Coordination
IEEE	Institute of Electrical and Electronics Engineers
IP	Internet Protocol
IFFT	Inverse Fast Fourier Transform
IPN	Interference Plus Noise
ILP	Integer Linear Programming
IMT	International Mobile Telecommunications
InH	Indoor Hotspot
IoT	Interference over Thermal

IR	Incremental Redundancy
IRC	Interference Rejection Combining
ISD	Inter-Site Distance
ISI	Intersymbol Interference
ITU	International Telecommunication Union
ITU-R	ITU Radiocommunication Sector
JT	Joint Transmission
KPI	Key Performance Indicator
LENA	LTE-EPC Network Simulator
LLR	Log-Likelihood Ratio
LoS	Line-of-Sight
LTE	Long Term Evolution
LMMSE	Linear Minimum Mean Squared Error
MAC	Medium Access Control
MCS	Modulation and Coding Scheme
MIESM	Mutual Information Effective SINR Mapping
MIMO	Multiple Input Multiple Output
MISO	Multiple Input Single Output
MLD	Maximum Likelihood Detection
MMIB	Mean Mutual Information per Bit
MMSE	Minimum Mean Squared Error
MRC	Maximum Ratio Combining
MS	Mobile Station
MU-MIMO	Multi-User MIMO
NACK	Negative Acknowledgement
NAICS	Network-Assisted Interference Cancellation and Suppression
NLoS	Non-Line-of-Sight
NOMA	Non-Orthogonal Multiple Access
O2I	Outdoor-to-Indoor
O2V	Outdoor-to-Vehicle
OI	Overload Indicator
OCC	Orthogonal Cover Code
OFDM	Orthogonal Frequency-Division Multiplexing
OFDMA	Orthogonal Frequency-Division Multiple Access
OLLA	Outer-Loop Link Adaptation
OLSM	Open-Loop Spatial Multiplexing
PBCH	Physical Broadcast Channel
PCFICH	Physical Control Format Indicator Channel
PDCCH	Physical Downlink Control Channel
PDCP	Packet Data Convergence Protocol
PDP	Power Delay Profile
PDF	Probability Density Function
PDSCH	Physical Downlink Shared Channel
PDU	Packet Data Unit
PF	Proportional Fair

P-GW	Packet Data Network Gateway
PHICH	Physical Hybrid-ARQ Indicator Channel
PMI	Precoding-Matrix Indicator
PRB	Physical Resource Block
PSS	Primary Synchronization Signal
ProSe	Proximity Services
PU²RC	Per-User Unitary Rate Control
PUCCH	Physical Uplink Control Channel
PUSCH	Physical Uplink Shared Channel
QAM	Quadrature Amplitude Modulation
QoS	Quality of Service
QPSK	Quadrature Phase Shift Keying
RAN1	Radio Access Network Working Group 1
RE	Resource Element
RF	Radio Frequency
RI	Rank Indicator
RLC	Radio Link Control
RMa	Rural Macro
RMS	Root Mean Square
RNTP	Relative Narrow Band Transmit Power
ROHC	Robust Header Compression
RR	Round-Robin
RRC	Radio Resource Control
RRM	Radio Resource Management
RSRP	Reference Signal Received Power
RTT	Round Trip Time
RTx	Retransmission
RTP	Real Time Transport
RV	Redundancy Version
Rx	Reception
SAR	Segmentation and Reassembly
SC-FDMA	Single-Carrier FDMA
SCM	Spatial Channel Model
SDMA	Space Division Multiple Access
SFR	Soft Frequency Reuse
SIC	Successive Interference Cancellation
SIMO	Single Input Multiple Output
SINR	Signal to Interference plus Noise Ratio
SIR	Signal to Interference Ratio
SISO	Single Input Single Output
SMa	Suburban Macro
SNR	Signal to Noise Ratio
SRS	Sounding Reference Signal
SSS	Secondary Synchronization Signal
SFBC	Space-Frequency Block Code

STBC	Space-Time Block Code
SVD	Singular Value Decomposition
TDD	Time-Division Duplexing
TDMA	Time-Division Multiple Access
TM	Transmission Mode
TM	Transparent Mode
TP	Transmission Point
TR	Technical Report
TS	Technical Specification
TTI	Transmission Time Interval
Tx	Transmission
UDP	User Datagram Protocol
UE	User Equipment
UL	Uplink
ULA	Uniform Linear Array
UM	Unacknowledged Mode
UMa	Urban Macro
UMTS	Universal Mobile Telecommunication System
UMi	Urban Micro
VoIP	Voice over IP
VoLTE	Voice over LTE
WINNER	Wireless World Initiative New Radio
WLAN	Wireless Local Area Network
WRC	World Radiocommunication Conference
ZF	Zero-Forcing

Nomenclature

x	italic lower- or uppercase letters denote scalar variables
$ x $	absolute value of a scalar x , i.e., $ x = \sqrt{x \cdot x^*}$
$\Im\{x\}$	imaginary part of a complex scalar x
$\Re\{x\}$	real part of a complex scalar x
x^*	complex conjugate of x
$\lfloor x \rfloor$	floor function, i.e., $\lfloor x \rfloor$ is the largest integer not greater than x
\mathbf{x}	italic lowercase boldface letters denote vector-valued variables; all vectors are assumed to be column vectors
x_i	i -th element of vector \mathbf{x}
$\ \mathbf{x}\ $	Euclidean norm of a vector \mathbf{x} , i.e., $\ \mathbf{x}\ = \sqrt{\sum_i x_i ^2}$
\mathbf{A}	italic uppercase boldface letters denote matrices
$[\mathbf{A}]_{i,j}$	element in row i and column j of matrix \mathbf{A}
$\mathbf{A}[xyz]$	submatrix resulting from matrix \mathbf{A} by selecting columns x, y, z, \dots , e.g., $\mathbf{A}[12]$ selects first two columns from \mathbf{A}
\mathbf{A}^T	the transpose of a matrix \mathbf{A} (also applies to vectors)
\mathbf{A}^H	the Hermitian transpose (i.e., the complex conjugate of the transpose) of a matrix \mathbf{A} (also applies to vectors)
\mathbf{I}_N	the $N \times N$ identity matrix. Also written as \mathbf{I} if the dimensions are clear from the context
$\ \mathbf{A}\ _F$	Frobenius norm of the matrix \mathbf{A} , i.e., $\ \mathbf{A}\ _F = \sqrt{\sum_i \sum_j [A]_{i,j} ^2}$
\mathbf{A}^+	Moore-Penrose pseudoinverse of a matrix \mathbf{A}
$\text{diag}(x_1, \dots, x_n)$	a diagonal matrix with the variables x_i on the diagonal. Also written as $\text{diag}(x_i)$ if the dimensions are clear from the context
$\text{diag}(\mathbf{A})$	yields the diagonal entries of matrix \mathbf{A} so that $\text{diag}(\text{diag}(\mathbf{A}))$ is a diagonal matrix where the diagonal entries are identical to those of matrix \mathbf{A}
$\text{rank}(\mathbf{A})$	the rank of matrix \mathbf{A}
$\text{tr}(\mathbf{A})$	trace of a quadratic matrix \mathbf{A}
$\mathbf{0}$	an all-zero vector or matrix
$\hat{a}, \hat{\mathbf{a}}, \hat{\mathbf{A}}$	a hat denotes some estimate of scalar, vector, or matrix variable
$\tilde{a}, \tilde{\mathbf{a}}, \tilde{\mathbf{A}}$	a tilde denotes some processed scalar, vector, or matrix variable (e.g., after filtering or cancellation)
\mathcal{S}	sets are denoted by calligraphic uppercase letters, e.g., the set \mathcal{S}
$ \mathcal{S} $	the cardinality of the set \mathcal{S}
$\mathcal{N}(\mu, \sigma^2)$	normal (Gaussian) distribution with parameters μ (mean) and σ^2 (variance)
$\mathcal{CN}(\mu, \sigma^2)$	circularly symmetric complex normal (Gaussian) distribution (of a random scalar) with mean μ and variance σ^2

$\mathcal{CN}(\boldsymbol{\mu}, \mathbf{R})$	circularly symmetric complex normal (Gaussian) distribution (of a random vector) with mean vector $\boldsymbol{\mu}$ and covariance matrix \mathbf{R}
$\mathcal{U}(a, b)$	uniform distribution on the support $[a, b]$
$\chi^2(k)$	Chi-squared distribution with k degrees of freedom
\star	convolution operator
$E[\cdot]$	expectation of some random variable
$\Pr(\cdot)$	probability of something
α	factor for fractional pathloss compensation (LTE UL power control)
α	smoothing factor for exponential smoothing used to compute past average throughput
β	ratio of received versus transmitted power
β	history exponent used in proportional fair metric
Γ	normal form of a non-cooperative game
$\bar{\gamma}_k$	average (past) throughput for user k (e.g., used in the PF metric)
γ	pathloss exponent
γ_i	normalized user throughput for user i
$\Delta(n, k)$	link adaptation offset for user k in TTI n
$\Delta\tau_{\max}$	channel length (i.e., maximum excess delay)
Δ_{ACK}	incremental (positive) SINR offset for each ACK used in outer-loop link adaptation
Δ_{NACK}	incremental (negative) SINR offset for each NACK used in outer-loop link adaptation
Δ_{SINR}	SINR gain, e.g., due to interference coordination scheme
$\delta(\cdot)$	Dirac delta function (unit impulse function)
η	system spectral efficiency
θ	elevation angle (up from the xy-plane)
$\theta_{3\text{ dB}}$	elevation pattern 3 dB beamwidth
θ_{tilt}	base station antenna downtilt angle
κ	cross-polarization power ratio
λ	eigenvalue
λ	wavelength
λ_0	wavelength at system center frequency
μ	mean (of Gaussian distribution)
$\boldsymbol{\zeta}^{(x,y)}$	vector holding large-scale parameter values for some point (x, y)
ζ	antenna element slant angle
ρ	ratio of users assigned to exclusive zone in FFR and SFR
$\boldsymbol{\Sigma}$	diagonal matrix of singular values
σ	standard deviation of Gaussian distribution
σ_{ASA}	parameter (standard deviation) of the distribution of the angular spread (arrival angles, at the MS) large-scale parameter σ_{AoA} in the IMT-Advanced channel model
σ_{ASD}	parameter (standard deviation) of the distribution of the angular spread (departure angles, at the BS) large-scale parameter σ_{AoD} in the IMT-Advanced channel model
σ_{AoA}	angular spread (arrival angles, at the MS) large-scale parameter as used in the IMT-Advanced channel model

σ_{AoD}	angular spread (departure angles, at the BS) large-scale parameter as used in the IMT-Advanced channel model
σ_{D}	Doppler spread
σ_{DS}	parameter (standard deviation) of the distribution of the delay spread
σ_{K}	σ_{delay} large-scale parameter in the IMT-Advanced channel model parameter (standard deviation) of the distribution of the Rician K parameter in the IMT-Advanced channel model
σ_{SF}	parameter (standard deviation) of the distribution of the shadow fading parameter in the IMT-Advanced channel model
σ_{delay}	delay spread large-scale parameter as used in the IMT-Advanced channel model
σ_{τ}	delay spread
σ_{th}^2	thermal noise power
ζ	singular value
τ_i	excess time delay of multipath i
Φ	potential function
Φ	initial phase
ϕ	azimuth angle (in the xy-plane)
$\phi_{3\text{dB}}$	azimuth pattern 3 dB beamwidth
ϕ_{BS}	base station antenna array broadside direction azimuth angle
$\varphi_{n,m}$	angle of arrival for ray m of cluster n
$\varphi_{n,m}$	angle of departure for ray m of cluster n
φ_i	phase shift experienced on antenna i
\mathcal{X}_t	action set consisting of all possible actions of a player t in a game
ω	ratio of the total system frequency bandwidth available to reuse 1 users in FFR and SFR
$a_{u,s,n,m}$	time-invariant complex-valued coefficient used for efficient CIR computation
\mathbb{B}	set (finite field) containing only the numbers 0 and 1, i.e., $\mathbb{B} = \{0, 1\}$
\mathbf{B}	lower triangular matrix of random numbers used in Bartlett decomposition, see definition in (4.27)
B	frequency bandwidth
B_c	coherence bandwidth
BLER*	target block error ratio (BLER) for outer-loop link adaptation
$b_{n,m}$	time-invariant real-valued coefficient used for efficient CIR computation
\mathbb{C}	set of all complex numbers
\mathcal{C}	codebook $\mathcal{C} = \{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_N\}$ with N vectors \mathbf{u}_n as entries
C	capacity
$c(t)$	channel impulse response at time t
$c_{u,s}$	channel impulse response between antennas u and s
c_i	Chi-squared distributed random coefficients used for Bartlett matrix decomposition
c_i	complex path coefficient (in channel impulse response)
\mathcal{D}	set of co-scheduled users (i.e., using same time/frequency resource)
$\mathbf{D}_M(i)$	matrix used as a factor in precoder calculation on subcarrier i with

	LTE open-loop spatial multiplexing using M layers
D_i	number of resources demanded in cell i
d	geographical distance
d_{BP}	breakpoint distance
d_{corr}	(de)correlation distance
d_I	distance to interfering base station
d_{in}	indoor distance
d_{out}	outdoor distance
d_S	distance to serving base station
d_s	distance between reference element and BS antenna array element number s
d_u	distance between reference element and MS antenna array element number u
\dim_x	the size of the x-dimension of, for example, a matrix
\dim_y	the size of the y-dimension of, for example, a matrix
E	error matrix
e	Euler's number, base of the natural logarithm
F	matrix of Fourier transform coefficients
F	electrical field vector
$F^{(H)}$	(electrical) field amplitude with horizontal polarization
$F^{(V)}$	(electrical) field amplitude with vertical polarization
F	number of frequency bins
f_c	system center/carrier frequency
f_i	subcarrier index
f_{idx}	frequency bin index
G	two-dimensional grid of real values: $G \in \mathbb{R}^{\dim_x \times \dim_y}$
G	3D antenna gain pattern
$G_{Azimuth}$	azimuth (horizontal) antenna gain pattern
$G_{Elevation}$	elevation (vertical) antenna gain pattern
G_{Min}	minimum antenna pattern gain (maximum attenuation)
G_{est}	processing gain for reference symbol based channel estimation
$g(\cdot)$	FIR filter impulse response used to introduce geographic correlations of large-scale channel model parameters
$g(\cdot)$	invertible function used for effective SINR mapping
H	MIMO frequency-domain channel matrix $H \in \mathbb{C}^{N_{Rx} \times N_{Tx}}$
H_{tk}	MIMO channel matrix between base station t and mobile station k
\bar{H}	equivalent virtual channel matrix for SFBC modeling
h	SIMO channel vector (h^T is the MISO channel vector)
h	complex-valued scalar channel transfer function
h_{BS}	base station antenna height
h_{MS}	mobile station antenna height
I	identity matrix
$I_{k',m}^{(t \rightarrow k)}$	interference power contribution $I_{k',m}^{(t \rightarrow k)}$ to the desired layer m at user k' as the sum of the contributions from all interfering layers $i = 1 \dots M_k$ that t transmits to user k

j	imaginary unit with $j^2 = -1$
\mathcal{K}	set of users (mobile stations)
K	Rician K factor (power ratio between LoS ray and all multi path components)
k	user (mobile station) index
L	lower triangular matrix (Cholesky factor)
L	control region length (number of symbols)
M	number of layers in spatial multiplexing transmission
M	number of rays per cluster
m	ray index (inside reflection cluster)
m	spatial layer index
$m_{t,t',f}$	binary variable replacing multiplication in integer linear program
MI	mutual information
\mathbb{N}	set of all natural numbers
N_{ACK}	number of HARQ ACKs signaled by some user during time span
N_{NACK}	number of HARQ NACKs signaled by some user during time span
N_{BS}	number of base stations
N_{MPC}	number of paths / reflection clusters in multipath propagation
N_{MS}	number of users (mobile stations)
$N_{\text{MU},t}$	number of co-scheduled users in MU-MIMO group being served on same time and frequency resource by BS t
N_{PRB}	number of PRBs
N_{Tx}	number of transmit antennas
N_{Rx}	number of receive antennas
N_{Sample}	number of samples
N_{SINR}	number of SINRs (e.g., for each subcarrier or PRB)
n	discrete time (TTI) index
n	multipath (cluster) index
N_{WallIn}	number of indoor apartment walls between femtocell and MS
N_{WallOut}	number of outdoor building walls between femtocell and MS
\mathbf{n}	noise vector
$\bar{\mathbf{n}}$	equivalent virtual noise vector for SFBC modeling
\mathbf{P}	precoder matrix
$\bar{\mathbf{P}}$	equivalent precoder for SFBC modeling (scaled identity)
\mathbf{p}	a rank-1 precoder (beamformer) with $\mathbf{p} \in \mathbb{C}^{N_{\text{Tx}}}$
P_0	uplink power control reference value
P_{Max}	maximum transmit power
P_n	fraction of power in path n
P_{tx}	transmit power
PL	pathloss
PL _{basic}	basic pathloss
PL _{in}	additional indoor pathloss
PL _{tw}	O2I wall penetration loss
PL _{veh}	vehicle penetration loss
PL _{WallIn}	apartment wall penetration loss in femtocell deployment
PL _{WallOut}	outdoor house wall penetration loss in femtocell deployment

\mathbf{Q}	transmit covariance matrix
\mathbb{R}	set of all real numbers
\mathbf{R}	correlation matrix
\mathbf{R}_{IPN}	interference plus noise covariance matrix
\mathbf{R}_n	noise covariance matrix
R	cell radius
R	reuse factor in frequency reuse scheme, e.g., $R = 3$ for reuse 3
r	receiver antenna index
$r(\cdot)$	maps a CQI value to the corresponding spectral efficiency, see Table B.1
\mathcal{S}	set of constellation symbols (modulation alphabet)
\mathbf{s}	vector of transmit symbols for spatial multiplexing transmission
\mathbf{s}	vector of correlated random variables used for large-scale parameter generation
$\bar{\mathbf{s}}$	SFBC transmit vector
S	number of BS antennas
S	rank of MIMO channel (matrix)
s	BS antenna index
s	transmitter antenna index
s_m	symbol on layer m before precoding
\hat{s}	receiver output (estimate) of transmitted symbol s
\mathcal{T}	set of transmitters (base stations)
T	number of base stations
T_c	coherence time
T_{CP}	cyclic prefix length (duration)
T_s	basic time unit in LTE, $T_s = 1 / (15\,000 \cdot 2048)$ s
T_U	useful symbol duration (i.e., without cyclic prefix)
t	time
t	base station index
$t(k)$	mapping from a mobile station index k to its serving base station with index t
\mathbf{U}	unitary matrix holding left-singular vectors in singular value decomposition
\mathbf{U}_M	DFT matrix used as a factor in precoder calculation with LTE open-loop spatial multiplexing using M layers
\mathbf{u}	unit norm vector element of codebook $\mathbf{u} \in \mathcal{C}$
U	number of MS antennas
U_t	utility function of player t
u	MS antenna index
\mathbf{V}	unitary matrix holding right-singular vectors in singular value decomposition
v	speed of movement
$v_{n,m}$	Doppler frequency component of ray m belonging to path n
\mathbf{W}	receive filter matrix
$W_p(\mathbf{R}, M)$	Wishart distribution with M degrees of freedom and covariance matrix \mathbf{R} of size $p \times p$

\mathbf{x}	vector of transmit symbols in complex-baseband representation
x	complex-baseband transmit signal
x_I	in-phase component of complex baseband signal x
x_Q	quadrature component of complex baseband signal x
x_{RF}	radio frequency transmit signal
$x(k, n, f_{\text{idx}})$	a mapping $x : \mathcal{K} \times \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{B}$ indicating if at TTI n user k is scheduled in PRB f_{idx} ($x(k, n, f_{\text{idx}}) = 1$) or not ($x(k, n, f_{\text{idx}}) = 0$)
\mathbf{x}_{-t}	tuple of actions by all players except for player t
\mathbf{y}	vector of received symbols in complex-baseband representation
$\bar{\mathbf{y}}$	equivalent virtual received vector for SFBC modeling
y	received signal in complex baseband representation
y_{RF}	radio frequency receive signal
\mathbf{Z}	matrix of channel impulse response coefficient for DFT by matrix multiplication

List of Figures

1.1.	Overview of 3GPP cellular standards and their peak data rates in the context of the IMT-2000 and IMT-Advanced classification [M.1645; Sch08]	2
1.2.	Exemplary cellular system with 7 base station sites serving 21 cell sectors	3
1.3.	Cell spectral efficiency distribution example	4
1.4.	High-level overview of determinants affecting system-level data rates .	5
1.5.	The chapters and the three pillars of the thesis	7
2.1.	Classification of MIMO schemes	12
2.2.	MIMO system model for a single link with precoder \mathbf{P} and Rx filter \mathbf{W}^H	16
2.3.	Downlink and uplink intra-cell and inter-cell interference situation (full interference only shown for cell 4)	23
2.4.	LTE protocol stack layers 1 and 2 with downlink data flow for one user .	25
2.5.	Structure of time and frequency resources for scheduling [36.211]	26
2.6.	Modulation symbol alphabets as defined in the LTE standard [36.211] . .	28
2.7.	Symbol durations and subcarrier spacing in the LTE resource grid	29
2.8.	LTE downlink physical layer processing for 2 users on 4 layers	29
2.9.	Different scheduling strategies in highly frequency-selective conditions .	39
2.10.	High-level overview of determinants affecting system spectral efficiency	41
2.11.	Selected determinants affecting link spectral efficiency	42
2.12.	Comparison of exemplary reference signals in LTE Rel-8 and Rel-10 . . .	44
2.13.	Selected intra-cell and inter-cell RRM aspects affecting spectral efficiency	47
2.14.	Simplified view of downlink transmission through LTE protocol stacks and over MIMO channel showing the scope of system-level simulations	49
2.15.	Simulation flow in <i>IMT_{Aphy}</i> with multi-threading of selected functions	56
3.1.	Overview of the main BS and user characteristics covered in this chapter	59
3.2.	Wraparound configuration surrounding a 21 cells (7 sites) scenario with 6 shifted copies of itself; arrows mark the offset of each copy	61
3.3.	Polar plots of antenna gain patterns	63
3.4.	Antenna setup for calibration simulations defined in TR [36.814]	64
3.5.	BS antenna configurations A) to E) with four antenna elements for IMT-Advanced evaluation simulations as defined in TR [36.814]	65
3.6.	Overview of propagation aspects covered by the IMT-A channel models	66
3.7.	Large-scale parameter filtering for the UMa scenario with LoS	71
3.8.	Large-scale calibration of <i>IMT_{Aphy}</i> against 3GPP results [R1-092019] . .	77
3.9.	Visualization of a small UMa scenario with 7 sites and 21 cells	78
3.10.	CDFs of the RMS delay spreads computed according to TR [25.996] . . .	79

3.11. CDFs of the RMS angular spreads computed according to TR [25.996] and reference curves [WIN10; Chi] in the UMa scenario	80
3.12. Comparing channel model implementations in the UMa scenario	81
4.1. Overview of link-to-system modeling methodology determining block error ratio for PDSCH allocation of single transport block (blue) on two PRBs and two layers [Iku13; BAS+05]	87
4.2. Receiver k' employing a filter matrix $\mathbf{W}_{k'}^H$ to receive two spatial layers $s_{k',1}$ and $s_{k',2}$ over a 2×2 MIMO channel $\mathbf{H}_{t'k'}$ from a desired transmitter t' using a precoding matrix $\mathbf{P}_{k'}$ with interference coming from an undesired transmitter t using a precoder \mathbf{P}_k to transmit to some other user $k \neq k'$	89
4.3. QPSK MMIB-based effective SINR computation example [EMD]	99
4.4. Block error ratios for different MCSs and block lengths based on LTE CQIs	101
4.5. Performance of different MCSs based on the 15 LTE CQIs	102
4.6. Comparison of the CDFs of the average per-user post-receiver SINRs and normalized throughput for the IMT-A scenarios	106
4.7. CDFs of the avg. per-user post-receiver SINRs and normalized throughput in the RMa scenario compared to 3GPP RAN1 results [R1-092742]	107
5.1. System-wide user throughput distribution achieved by RR (dashed line) and PF scheduling (solid lines) depending on PF history exponents β	112
5.2. Comparison of cell spectral efficiency and post-scheduling average SINRs per user for RR (dashed lines) and PF (solid lines) schedulers in the UMi scenario; PF schedulers with exponents β between 0 and 5	113
5.3. Scheduling gains in cell spectral efficiency, cell-edge user throughput, and post-scheduling average SINR (right-hand axis) for RR and PF schedulers versus subband size (error bars indicate 95% confidence intervals)	115
5.4. Multi-user diversity gains achieved by PF and RR schedulers: cell spectral efficiency versus average number of users per cell in the UMi scenario	117
5.5. Scheduling gains and coherence time vs. user speed in the UMi scenario	120
5.6. Link adaptation mismatches (actual – estimated SINR), resulting BLERs, and spectral efficiencies for RR and PF schedulers with and without OLLA in the UMi and InH scenarios versus user velocities (legends for Figs. (b) and (c) see Fig. (a))	123
5.7. Comparison of link adaptation mismatch between RR and PF scheduling with and without OLLA for 10 km/h and 30 km/h	125
5.8. Throughput (TP) gains in terms of system spectral efficiency (top left) and cell-edge user throughput (bottom left) depending on the number of Rx antennas, receive filter, and link adaptation scheme in the UMi scenario; post-receiver SINRs separated into mean SNR and mean IoT (on the right)	127
5.9. Normalized polar plots in (linear scale) of 4 Rx antennas rank-1 receive filters for different IPN covariance matrix estimation schemes	128

5.10. Instantaneous SINR distribution for different serving channel estimation processing gains G_{est} compared to perfect serving channel estimation for 2 and 4 Rx antennas with MRC and IRC receivers in the UMi scenario	130
5.11. Instantaneous SINR achieved by the MMSE(-IRC) receiver using different IPN covariance estimation methods in the UMa scenario	132
5.12. Cell spectral efficiency and normalized cell-edge user throughput achieved by MMSE(-IRC) linear receivers with different IPN covariance matrix estimation models in the UMa scenario	133
5.13. Comparison of instantaneous SINR and per-user normalized throughput between single BS Tx antenna (TM1) and 2 Tx antenna SFBC transmit diversity operation (TM2) with special 1 PRB per user round-robin scheduling for 1, 2, and 4 Rx antennas at the MS	134
5.14. System-level performance of SFBC with channel-adaptive scheduling . .	136
5.15. Polar plots in linear scale of the rank-1 beamforming patterns from the LTE PMI codebook (see Tables B.3 to B.6 and Table B.7 in the Appendix) after application of the IMT-A azimuth sector pattern (see (3.3) and Fig. 3.3) with a boresight direction of 0° . Antennas are symbolized by black dots; see Fig. 5.16 for the PMI color code.	139
5.16. Geographic distribution of PMIs in the RMa scenario (high LoS prob.) .	141
5.17. Performance of TM6 in the RMa scenario	144
5.18. Illustration of the flashlight effect	146
5.19. Link adaptation mismatch and per-cell system spectral efficiency distributions demonstrating the flashlight effect in the RMa scenario	147
5.20. Impact of 4×4 MIMO antenna configurations on the rank mixes and per-user throughput distribution in the RMa scenario; four antenna array combinations (BS \leftrightarrow MS, see Fig. 3.5 on page 65 as a reference), each for TM3 and TM4 as well as TM1 as a reference for the throughput . . .	153
5.21. Impact of 4×4 MIMO antenna configurations on the rank mixes and per-user throughput distribution in the InH scenario; four antenna array combinations (BS \leftrightarrow MS, see Fig. 3.5 on page 65 as a reference), each for TM3 and TM4 as well as TM1 as a reference for the throughput	156
5.22. System and cell-edge user spectral efficiency in the UMi scenario with a 2×2 antenna configuration (E-E), evaluating those 10 % of users with TM and max. rank according to the y-axis while the other 90 % of users cause interference with TMs and max. ranks as indicated on the x-axis .	160
5.23. System and cell-edge user spectral efficiency in the InH scenario with a 4×4 antenna configuration (E-E) with an MMSE-MRC receiver, evaluating those 10 % of users with TM and max. rank according to the y-axis while the other 90 % of users cause interference with TMs and max. ranks as indicated on the x-axis	162

5.24. System and cell-edge user spectral efficiency in the InH scenario with a 4×4 antenna configuration (E-E) with an MMSE-IRC receiver, evaluating those 10 % of users with TM and max. rank according to the y-axis while the other 90 % of users cause interference with TMs and max. ranks as indicated on the x-axis	163
5.25. Geometry of simplified cellular model with SIR distribution for $\gamma = 4$.	164
5.26. Cell spectral efficiency and area spectral efficiency over ISD in the down-link of a 1×2 UMa deployment with MRC receivers	166
6.1. Classification of single- and multi-cell RRM methods w.r.t. coordination scope, update rate, and signaling load (low: light blue, high: dark blue) .	169
6.2. Overview of static frequency reuse schemes	171
6.3. Required SINR improvement to achieve equal mutual information rate with reuse 3 compared to reuse 1 depending on long-term SINR	173
6.4. Geographical distribution of the SINR and MI gain in the UMa scenario comparing reuse 3 to reuse 1 with a SISO configuration and RR scheduling	174
6.5. Distribution of the per-user reuse 3 versus reuse 1 MI gain according to equation (2.27) in the UMi scenario depending on antenna configuration and scheduler; dots indicate scenario-wide average of reuse 3 gain . . .	175
6.6. Cell-edge vs. system-spectral efficiency for FFR (user ratio $\rho = 0.2 \dots 0.6$), reuse 1, and reuse 3 with different MIMO schemes and RR and PF schedulers (reuse 1 with PF history exponents $\beta = 0.5 \dots 2.5$); see also Table 6.1	177
6.7. Dual strip model with two single-floor houses [36.814] where in 20 out of 40 apartments a femtocell is active	181
6.8. Convergence to Nash equilibria and comparison of equilibrium efficiencies to benchmarks for the sequential and p -persistent schemes	186
6.9. Exemplary beamformer (PMI) to subband mapping with 4 beams per cell	188
6.10. Polar plot of the per-layer and total radiated power gain using 1–4 columns of the unitary PMI 0 precoding matrix (see Table B.3) with a linear array of four $\lambda/2$ -spaced antennas	193
6.11. Distribution of SINR per user and layer for different MU-MIMO and reference schemes in a UMi scenario with on average 25 users per cell .	197
6.12. Distribution of SINR estimation mismatch for different MU-MIMO and reference schemes in a UMi scenario with on average 25 users per cell .	198
6.13. Distribution of normalized user throughputs for different MU-MIMO and reference schemes in a UMi scenario with on average 25 users per cell	199
6.14. Impact of multi-user diversity on average SINRs, idealized SINR estimation mismatches, and consequently throughput with presented MU-MIMO schemes with TM4 and TM6 as a reference	201
A.1. Throughput (TP) gains in terms of system spectral efficiency (top left) and cell-edge user throughput (bottom left) depending on the number of Rx antennas and the receive filter and link adaptation scheme in the UMa scenario with the underlying post-combining mean per-user (linear) average SINRs separated into mean SNR and mean IoT	213

A.2. Instantaneous SINR achieved by the 4 Rx MMSE(-IRC) receiver using different IPN covariance estimation methods in the UMi scenario	214
A.3. Cell spectral efficiency and normalized cell-edge user throughput achieved by MMSE linear receiver with different IPN covariance matrix estimation models in the UMi scenario. See Table 5.9 for simulation parameters.	215
A.4. Geographic distribution of PMIs in the UMa scenario (low LoS prob.) . .	216
A.5. Performance of TM6 in the UMa scenario	217
A.6. Impact of 4×4 MIMO antenna configurations on the rank mixes and per-user throughput distribution in the UMa scenario; four antenna array combinations (BS \leftrightarrow MS), each for TM3 and TM4 as well as TM1 as a reference for the throughput	218
A.7. Impact of 4×4 MIMO antenna configurations on the rank mixes and per-user throughput distribution in the UMi scenario; four antenna array combinations (BS \leftrightarrow MS), each for TM3 and TM4 and TM1 as a reference for the throughput	219
A.8. System and cell-edge user spectral efficiency in the InH scenario with a 4×4 antenna configuration (E-E), evaluating those 10 % of users with TM and max. rank according to the y-axis while the other 90 % of users cause interference with TMs and max. ranks as indicated on the x-axis .	220

List of Tables

2.1.	LTE-Advanced DL transmission modes in LTE Rel-11 [36.213, Sec. 7.1]	34
2.2.	Exemplary transceiver performance with net rate and attribution of individual losses compared to capacity at 10.4 dB according to [RCMW11]	43
2.3.	PHY system design overhead for different number of antenna ports (with 12×14 REs, normal cyclic prefix, 10 MHz system bandwidth)	45
3.1.	Overview of deployment scenarios and simulation parameters [M.2135]	60
3.2.	Common deployment scenario parameters [M.2135]	60
3.3.	Line-of-sight propagation probability model and simulation results	68
3.4.	UMa large-scale parameter distribution parameters [M.2135]	70
3.5.	Spatial channel model parameters	74
4.1.	3GPP assumptions for IMT-Advanced simulator calibration [36.814]	104
5.1.	General simulation parameters for the simulations in this chapter	110
5.2.	Specific simulation parameters for Figs. 5.1 and 5.2	112
5.3.	Specific simulation parameters for Fig. 5.3	114
5.4.	Specific simulation parameters for Fig. 5.4	116
5.5.	Specific simulation parameters for Fig. 5.5	119
5.6.	Specific simulation parameters for Figs. 5.6 and 5.7	123
5.7.	Specific simulation parameters for Fig. 5.8 and Fig. A.1	126
5.8.	Specific simulation parameters for Fig. 5.10	130
5.9.	Specific simulation parameters for Figs. 5.11, 5.12, A.2 and A.3	131
5.10.	Specific simulation parameters for Fig. 5.13 and Table 5.11	134
5.11.	Initial transmission BLER performance of TM1 and TM2 with special 1 PRB per user round-robin scheduling	135
5.12.	Specific simulation parameters for Fig. 5.14	136
5.13.	Simulation parameters for Figs. 5.17 and A.5 and Tables 5.14 and 5.15	142
5.14.	TM6 system spectral efficiencies in bit/s/Hz depending on the receiver type, antenna configuration, and antenna spacing in the RMa scenario	143
5.15.	TM6 cell-edge user throughput in bit/s/Hz depending on the receiver type, antenna configuration, and antenna spacing in the RMa scenario	145
5.16.	Simulation parameters for the simulations presented in Section 5.5.2	151
5.17.	PHY losses in the UMi and InH scenario determined in system-level simulations as $1 - \frac{B \cdot \gamma}{M}$ for different MIMO configurations	158
5.18.	Simulation parameters for the simulations presented in Section 5.6	159
6.1.	Simulation parameters for Figs. 6.5 and 6.6	176
6.2.	Overview of femtocell deployment and pathloss model parameters	184

6.3.	Efficiency of Nash equilibria compared to benchmarks after 250 iterations	186
6.4.	Simulation parameters for the simulations presented in Section 6.2.3	197
6.5.	Cell spectral efficiency and cell-edge user throughput in bit/s/Hz in the UMi scenario for MU-MIMO and reference schemes with 25 users per cell	200
A.1.	Frequency of PMI selection, 2 Tx antennas	214
A.2.	Frequency of PMI selection, 4 Tx antennas	214
A.3.	TM6 system spectral efficiencies in bit/s/Hz depending on the receiver type, antenna configuration, and antenna spacing in the UMa scenario	214
A.4.	TM6 cell-edge user throughput in bit/s/Hz depending on the receiver type, antenna configuration, and antenna spacing in the UMa scenario	215
A.5.	System cell-edge user spectral efficiency of TM3 and TM4 compared to TM1 depending on scenario and antenna array configurations	215
A.6.	Overview of system spectral efficiency and cell-edge user throughput for open- and closed-loop MIMO compared to TM1 depending on scenario and antenna array configurations	216
B.1.	LTE CQI definition [36.213, Sec. 7.2.3]	221
B.2.	LTE open-loop spatial multiplexing precoding factors	222
B.3.	LTE 4 Tx Codebook, PMIs 0–3 (column 1: blue, column 2: red, column 3: green, column 4: cyan, total: black)	223
B.4.	LTE 4 Tx Codebook, PMIs 4–7 (column 1: blue, column 2: red, column 3: green, column 4: cyan, total: black)	224
B.5.	LTE 4 Tx Codebook, PMIs 8–11 (column 1: blue, column 2: red, column 3: green, column 4: cyan, total: black)	225
B.6.	LTE 4 Tx Codebook, PMIs 12–15 (column 1: blue, column 2: red, column 3: green, column 4: cyan, total: black)	226
B.7.	LTE 2 Tx Codebook, PMIs 0–3 (column 1: blue, column 2: red, total: black)	227

Publications and Presentations by the Author

- [Ell07a] R. Pabst, J. Ellenbeck, M. Schinnenburg, and C. Hoymann, "System level performance of cellular WiMAX IEEE 802.16 with SDMA-enhanced medium access," in *Proc. IEEE Wireless Communications & Networking Conference (WCNC 2007)*, Hong Kong, China, Mar. 2007.
- [Ell07b] C. Hoymann, J. Ellenbeck, R. Pabst, and M. Schinnenburg, "Evaluation of grouping strategies for an hierarchical SDMA/TDMA scheduling process," in *Proc. IEEE International Conference on Communications (ICC 2007)*, Glasgow, UK, Jun. 2007.
- [Ell08] J. Ellenbeck, C. Hartmann, and L. Berlemann, "Decentralized inter-cell interference coordination by autonomous spectral reuse decisions," in *Proc. European Wireless Conference (EW 2008)*, Prague, Czech Republic, Jun. 2008.
- [Ell09a] J. Ellenbeck, H. Al-Shatri, and C. Hartmann, "Performance of decentralized interference coordination in the LTE uplink," in *Proc. IEEE Vehicular Technology Conference (VTC-Fall 2009)*, Anchorage, USA, Sep. 2009.
- [Ell09b] J. Ellenbeck, J. Schmidt, U. Korger, and C. Hartmann, "A concept for efficient system-level simulations of OFDMA systems with proportional fair fast scheduling," in *Proc. 5th IEEE Broadband Wireless Access Workshop co-located with IEEE Globecom 2009*, Honolulu, Hawaii, USA, Nov. 2009.
- [Ell10] J. Ellenbeck, M. Hammoud, B. Lazarov, and C. Hartmann, "Autonomous beam coordination for the downlink of an IMT-Advanced cellular system," in *Proc. European Wireless Conference (EW 2010)*, Lucca, Italy, Apr. 2010.
- [Ell11] J. Ellenbeck, O. Mushtaq, F. Sheikh, and S. Q. Shahrukh, "Calibrating an efficient C++ implementation of the ITU-R M.2135 channel model for use in system-level simulations," in *Proc. International Workshop on Propagation and Channel Modeling for Next-Generation Wireless Networks (IWPCM 2011)*, Lyon, France, Mar. 2011.

- [Ell12a] J. Ellenbeck, "System-level simulations with the IMT-Advanced channel model," in *LTE-Advanced and Next Generation Wireless Networks: Channel Modelling and Propagation*, G. de la Roche, A. Alayon-Glazunov, and B. Allen, Eds., John Wiley & Sons, 2012, ch. 13, pp. 349–387.
- [Ell12b] J. Ellenbeck, A. Dotzler, and W. Utschick, "Interference robustness for MIMO networks — system-level performance evaluation for LTE-Advanced," in *Proc. International Workshop on Emerging Technologies for LTE-Advanced and Beyond-4G in conjunction with IEEE Globecom 2012*, Anaheim, California, USA, Dec. 2012.
- [EllPres09] C. Hartmann and J. Ellenbeck, "Decentralized interference coordination for LTE," presented at the 30. Treffen der VDE/ITG-Fachgruppe 5.2.4, Workshop on Cooperation and Self-Organization in Communication Networks, Klagenfurt, Austria, Jun. 29, 2009.
- [EllPres10a] J. Ellenbeck, "Interference management: From autonomous to closely coordinated approaches," presented at the 32. Treffen der VDE/ITG-Fachgruppe 5.2.4, Workshop on Interference Management and Cooperation Strategies in Communication Networks, Darmstadt, Germany, Feb. 18, 2010.
- [EllPres10b] —, "Self-organizing inter-cell interference coordination," presented at the VDE/ITG-Fachgruppe 5.2.1, Workshop Selbstorganisation, Stuttgart, Germany, Oct. 6, 2010.
- [EllPres11] —, "IMTAphy: An efficient M.2135 spatial channel model and link-to-system mapping implementation for system-level simulations of IMT-Advanced systems," presented at the 36. Treffen der VDE/ITG-Fachgruppe 5.2.4, Workshop on Simulation and Prototyping Environments for Mobile/Wireless Research, Aachen, Germany, Jul. 13, 2011.
- [EllWiWi07] —, "Game theoretic modeling and evaluation of the radio resource competition in future cellular wireless networks," Master's thesis, Rheinisch-Westfälische Technische Hochschule Aachen, Aachen, Germany, Sep. 2007.

General Literature

- [AG99] M.-S. Alouini and A. Goldsmith, "Area spectral efficiency of cellular mobile radio systems," *IEEE Transactions on Vehicular Technology*, vol. 48, no. 4, pp. 1047–1066, Jul. 1999.
- [Ala98] S. Alamouti, "A simple transmit diversity technique for wireless communications," *IEEE Journal on Selected Areas in Communications*, vol. 16, no. 8, pp. 1451–1458, Oct. 1998.
- [AvHK+07] K. Aardal, S. van Hoesel, A. Koster, C. Mannino, and A. Sassano, "Models and solution techniques for frequency assignment problems," *Annals of Operations Research*, vol. 153, no. 1, pp. 79–129, 2007.
- [AVWS12] R. Akl, S. Valentin, G. Wunder, and S. Stanczak, "Compensating for CQI aging by channel prediction: the LTE downlink," in *Proc. IEEE Global Communications Conference (GLOBECOM 2012)*, Anaheim, California, USA, Dec. 2012, pp. 4821–4827.
- [B B05] B. Badic, "Space-time block coding for multiple antenna systems," PhD thesis, Institut für Nachrichtentechnik und Hochfrequenztechnik, Technische Universität Wien, Vienna, Austria, 2005.
- [BA09] G. Bauch and T. Abe, "On the parameter choice for cyclic delay diversity based precoding with spatial multiplexing," in *Proc. IEEE Global Telecommunications Conference (GLOBECOM 2009)*, Honolulu, Hawaii, USA, Dec. 2009.
- [Bar33] M. Bartlett, "On the theory of statistical regression," *Proceedings of the Royal Society of Edinburgh*, vol. 53, pp. 260–283, 1933.
- [BAS+05] K. Brueninghaus, D. Astély, T. Sälzer, S. Visuri, A. Alexiou, S. Karger, and G. Seraji, "Link performance models for system level simulations of broadband radio access systems," in *Proc. IEEE Personal, Indoor and Mobile Radio Communications Conference (PIMRC 2005)*, Berlin, Germany, Sep. 2005, pp. 2306–2311.
- [BCG+07] E. Biglieri, A. Calderbank R. Constantinides, A. Goldsmith, A. Paulraj, and H. Poor, *MIMO Wireless Communications*. Cambridge University Press, 2007.

- [BCH+95] B. Boehm, B. Clark, E. Horowitz, C. Westland, R. Madachy, and R. Selby, "Cost models for future software life cycle processes: COCOMO 2.0," *Annals of Software Engineering*, vol. 1, no. 1, pp. 57–94, 1995.
- [BGW09] M. Bohge, J. Gross, and A. Wolisz, "Optimal power masking in soft frequency reuse based OFDMA networks," in *Proc. European Wireless Conference (EW 2009)*, Aalborg, Denmark, 2009, pp. 162–166.
- [BHS05] D. Baum, J. Hansen, and J. Salo, "An interim channel model for beyond-3G systems: extending the 3GPP spatial channel model (SCM)," in *Proc. IEEE Vehicular Technology Conference (VTC-Spring 2005)*, vol. 5, Stockholm, Sweden, May 2005, pp. 3132–3136.
- [BLS+13] A. Benjebbour, A. Li, Y. Saito, Y. Kishiyama, A. Harada, and T. Nakamura, "System-level performance of downlink NOMA for future LTE enhancements," in *Proc. IEEE Global Telecommunications Conference (GLOBECOM 2013)*, Atlanta, Georgia, USA, Dec. 2013.
- [BMKS09] D. Bültmann, M. Mühleisen, K. Klagges, and M. Schinnenburg, "openWNS — open wireless network simulator," in *Proc. European Wireless Conference (EW 2009)*, Aalborg, Denmark, May 2009, pp. 211–215.
- [BMRN11] N. Baldo, M. Miozzo, M. Requena-Esteso, and J. Nin-Guerrero, "An open source product-oriented LTE network simulator based on ns-3," in *Proc. ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM '11)*, Miami, Florida, USA, 2011, pp. 293–298.
- [Bon06] E. Bonek, "The MIMO radio channel," in *Proc. Nordic Antenna Symposium*, May 2006, pp. 23–34.
- [CAG08] V. Chandrasekhar, J. Andrews, and A. Gatherer, "Femtocell networks: a survey," *IEEE Communications Magazine*, vol. 46, no. 9, pp. 59–67, 2008.
- [CDF+11] Y. Chen, F. Dhiem, K. Fang, Y. Liu, S. Martinez Lopez, L. H. Nguyen, B. Ning, H. Schoeneich, A. Tyrrell, R. Visoz, and T. Wild, "Advanced link-to-system modeling," *Advanced Radio Interface Technologies for 4G Systems (ARTIST4G)*, Deliverable D2.3, Mar. 2011.
- [CG01] S. T. Chung and A. Goldsmith, "Degrees of freedom in adaptive modulation: a unified view," *IEEE Transactions on Communications*, vol. 49, no. 9, pp. 1561–1571, Sep. 2001.
- [CGE03] S. Catreux, L. Greenstein, and V. Erceg, "Some results and insights on the performance gains of MIMO systems," *IEEE Journal on Selected Areas in Communications*, vol. 21, no. 5, pp. 839–847, 2003.

- [CJ08] V. Cadambe and S. Jafar, "Interference alignment and degrees of freedom of the k-user interference channel," *IEEE Transactions on Information Theory*, vol. 54, no. 8, pp. 3425–3441, 2008.
- [Cos83] M. H. M. Costa, "Writing on dirty paper," *IEEE Transactions on Information Theory*, vol. 29, no. 3, pp. 439–441, 1983.
- [Cox12] C. Cox, *An introduction to LTE: LTE, LTE-Advanced, SAE and 4G mobile communications*. Wiley, 2012.
- [CPG+13] F. Capozzi, G. Piro, L. Grieco, G. Boggia, and P. Camarda, "Downlink packet scheduling in LTE cellular networks: key design issues and a survey," *IEEE Communications Surveys Tutorials*, vol. 15, no. 2, pp. 678–700, 2013.
- [CS03] G. Caire and S. Shamai, "On the achievable throughput of a multiantenna Gaussian broadcast channel," *IEEE Transactions on Information Theory*, vol. 49, no. 7, pp. 1691–1706, Jul. 2003.
- [CT06] T. Cover and J. Thomas, *Elements of Information Theory*. John Wiley and Sons, 2006.
- [CTB98] G. Caire, G. Taricco, and E. Biglieri, "Bit-interleaved coded modulation," *IEEE Transactions on Information Theory*, vol. 44, no. 3, pp. 927–946, 1998.
- [CWK+09] C.-C. Chong, F. Watanabe, K. Kitao, T. Imai, and H. Inamura, "Evolution trends of wireless MIMO channel modeling towards IMT-Advanced," *IEICE Transactions on Communications*, vol. 92, no. 9, pp. 2773–2788, Sep. 2009.
- [Da 12] G. W. O. Da Costa, "Dynamic spectrum sharing among femtocells — coping with spectrum scarcity in 4G and beyond," PhD thesis, Aalborg University, Aalborg, Denmark, Aug. 2012.
- [DB07] G. Dietl and G. Bauch, "Linear precoding in the downlink of limited feedback multiuser MIMO system," in *Proc. IEEE Global Telecommunications Conference (GLOBECOM 2007)*, Washington DC, USA, Nov. 2007.
- [dCCKM10] G. da Costa, A. Cattoni, I. Kovács, and P. Mogensen, "A scalable spectrum-sharing mechanism for local area network deployment," *IEEE Transactions on Vehicular Technology*, vol. 59, no. 4, pp. 1630–1645, May 2010.
- [dFS10] R. de Francisco and D. T. Slock, "An optimized unitary beamforming technique for MIMO broadcast channels," *IEEE Transactions on Wireless Communications*, vol. 9, no. 3, pp. 990–1000, Mar. 2010.
- [DHL+11] M. Dohler, R. Heath, A. Lozano, C. Papadidas, and R. Valenzuela, "Is the PHY layer dead?" *IEEE Communications Magazine*, vol. 49, no. 4, pp. 159–165, Apr. 2011.

- [DLU09] G. Dietl, O. Labrèche, and W. Utschick, "Channel vector quantization for multiuser MIMO systems aiming at maximum sum rate," in *Proc. IEEE Global Telecommunications Conference (GLOBECOM 2009)*, Honolulu, Hawaii, USA, Nov. 2009.
- [DPS11] E. Dahlman, S. Parkvall, and J. Sköld, *4G LTE / LTE-Advanced for mobile broadband*. Academic Press, 2011.
- [DPV07] J. Duplicy, D. P. Palomar, and L. Vandendorpe, "Adaptive orthogonal beamforming for the MIMO broadcast channel," in *Proc. IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMPSPAP 2007)*, St. Thomas, Virgin Islands, Dec. 2007, pp. 77–80.
- [DS05] G. Dimic and N. Sidiropoulos, "On downlink beamforming with greedy user selection: performance analysis and a simple new algorithm," *IEEE Transactions on Signal Processing*, vol. 53, no. 10, pp. 3857–3868, Oct. 2005.
- [Due07] A. Duel-Hallen, "Fading channel prediction for mobile radio adaptive transmission systems," *Proceedings of the IEEE*, vol. 95, no. 12, pp. 2299–2313, 2007.
- [EPH13] O. El Ayach, S. Peters, and R. Heath Jr., "The practical challenges of interference alignment," *IEEE Wireless Communications*, vol. 20, no. 1, pp. 35–42, 2013.
- [EVBH08] J. Eberspächer, H.-J. Vögel, C. Bettstetter, and C. Hartmann, *GSM — Architecture, Protocols and Services*. John Wiley & Sons, 2008.
- [FDM+12] G. Fodor, E. Dahlman, G. Mildh, S. Parkvall, N. Reider, G. Miklós, and Z. Turányi, "Design aspects of network assisted device-to-device communications," *IEEE Communications Magazine*, vol. 50, no. 3, pp. 170–177, 2012.
- [FG98] G. Foschini and M. Gans, "On limits of wireless communications in a fading environment when using multiple antennas," *Wireless Personal Communications*, vol. 6, no. 3, pp. 311–335, 1998.
- [FL98] D. Fudenberg and D. K. Levine, *The Theory of Learning in Games*. MIT Press, 1998.
- [FMMO99] A. Furuskär, S. Mazur, F. Müller, and H. Olofsson, "EDGE: enhanced data rates for GSM and TDMA/136 evolution," *IEEE Personal Communications*, vol. 6, no. 3, pp. 56–66, 1999.
- [FN94] C. Farsakh and J. Nossek, "Application of space division multiple access to mobile radio," in *Proc. Personal, Indoor and Mobile Radio Communications Conference (PIMRC) 1994*, The Hague, The Netherlands, Sep. 1994.
- [Gal08] R. Gallager, *Principles of Digital Communication*. Cambridge University Press, 2008.

- [GBGP02] D. Gesbert, H. Bölcskei, D. Gore, and A. Paulraj, "Outdoor MIMO wireless channels: models and performance prediction," *IEEE Transactions on Communications*, vol. 50, no. 12, pp. 1926–1934, 2002.
- [GKP+12] L. Garcia, I. Kovács, K. Pedersen, G. Costa, and P. Mogensen, "Autonomous component carrier selection for 4G femtocells — a fresh look at an old problem," *IEEE Journal on Selected Areas in Communications*, vol. 30, no. 3, pp. 525–537, Apr. 2012.
- [GML08] R. Giuliano, C. Monti, and P. Loretì, "WiMAX fractional frequency reuse for rural environments," *IEEE Wireless Communications*, vol. 15, no. 3, pp. 60–65, 2008.
- [God97] L. Godara, "Application of antenna arrays to mobile communications. II. beam-forming and direction-of-arrival considerations," *Proceedings of the IEEE*, vol. 85, no. 8, pp. 1195–1245, 1997.
- [Gol05] A. Goldsmith, *Wireless Communications*. Cambridge University Press, 2005.
- [Hal83] S. Halpern, "Reuse partitioning in cellular systems," in *Proc. IEEE Vehicular Technology Conference (VTC 1983)*, vol. 33, May 1983, pp. 322–327.
- [Hat77] D. Hatfield, "Measures of spectral efficiency in land mobile radio," *IEEE Transactions on Electromagnetic Compatibility*, vol. EMC-19, no. 3, pp. 266–268, 1977.
- [HCM+12] C. Hoymann, W. Chen, J. Montojo, A. Golitschek, C. Koutsimanis, and X. Shen, "Relaying operation in 3GPP LTE: challenges and solutions," *IEEE Communications Magazine*, vol. 50, no. 2, pp. 156–162, 2012.
- [HCPR12] H. Huh, G. Caire, H. Papadopoulos, and S. Ramprasad, "Achieving "massive MIMO"spectral efficiency with a not-so-large number of antennas," *IEEE Transactions on Wireless Communications*, vol. 11, no. 9, pp. 3226–3239, Sep. 2012.
- [HLKC13] C. Hoymann, D. Larsson, H. Koorapaty, and J.-F. Cheng, "A lean carrier for LTE," *IEEE Communications Magazine*, vol. 51, no. 2, pp. 74–80, 2013.
- [Hol00] J. Holtzman, "CDMA forward link waterfilling power control," in *Proc. Vehicular Technology Conference (VTC-Spring 2000)*, vol. 3, Tokyo, Japan, May 2000, pp. 1663–1667.
- [Hol01] —, "Asymptotic analysis of proportional fair algorithm," in *Proc. Personal, Indoor and Mobile Radio Communications Conference (PIMRC 2001)*, vol. 2, San Diego, California, USA, Sep. 2001, F-33–F-37.

- [Hol13] H. Holtkamp, "Enhancing the energy efficiency of radio base stations," PhD thesis, University of Edinburgh, Edinburgh, Scotland, Great Britain, Nov. 2013.
- [HPV12] H. Huang, C. Papadias, and S. Venkatesan, *MIMO Communication for Cellular Networks*. Springer, 2012.
- [HPWZ13] R. Heath, S. Peters, Y. Wang, and J. Zhang, "A current perspective on distributed antenna systems for the downlink of cellular systems," *IEEE Communications Magazine*, vol. 51, no. 4, pp. 161–167, 2013.
- [HT11] H. Holma and A. Toskala, *LTE for UMTS: Evolution to LTE-Advanced*. Wiley, 2011.
- [HvR09] P. Hosein and C. van Rensburg, "On the performance of downlink beamforming with synchronized beam cycles," in *Proc. IEEE Vehicular Technology Conference (VTC-Spring 2009)*, Barcelona, Spain, Apr. 2009.
- [IDM+11] R. Irmer, H. Droste, P. Marsch, M. Grieger, G. Fettweis, S. Brueck, H.-P. Mayer, L. Thiele, and V. Jungnickel, "Coordinated multipoint: concepts, performance, and field trial results," *IEEE Communications Magazine*, vol. 49, no. 2, pp. 102–111, 2011.
- [Iku13] J. C. Ikuno, "System level modeling and optimization of the LTE downlink," PhD thesis, Technische Universität Wien, Vienna, Austria, 2013.
- [IN07] M. Ivrlač and J. A. Nossek, "Intercell-interference in the Gaussian MISO broadcast channel," in *Proc. IEEE Global Telecommunications Conference (GLOBECOM 2007)*, Washington DC, USA, Nov. 2007.
- [Ivr05] M. Ivrlač, "Wireless MIMO systems — models, performance, optimization," PhD thesis, Technische Universität München, Munich, Germany, 2005.
- [IWR10] J. C. Ikuno, M. Wrulich, and M. Rupp, "System level simulation of LTE networks," in *Proc. IEEE Vehicular Technology Conference (VTC-Spring 2010)*, Taipei, Taiwan, May 2010.
- [Jak74] W. Jakes, *Microwave mobile communications*, ser. IEEE Press classic reissue. IEEE Press, 1974.
- [Jal06] J. Jaldén, "Detection for multiple input multiple output channels: analysis of sphere decoding and semidefinite relaxation," PhD thesis, KTH, Stockholm, Sweden, Dec. 2006.
- [JBTJ12] S. Jaeckel, K. Borner, L. Thiele, and V. Jungnickel, "A geometric polarization rotation model for the 3-D spatial channel model," *IEEE Transactions on Antennas and Propagation*, vol. 60, no. 12, pp. 5966–5977, 2012.

- [JCH84] R. K. Jain, D.-M. W. Chiu, and W. R. Hawe, "A quantitative measure of fairness and discrimination for resource allocation in shared computer system," Eastern Research Laboratory, Digital Equipment Corporation, Tech. Rep. DEC-TR-301, Sep. 1984.
- [JCK+11] M. Jain, J. I. Choi, T. Kim, D. Bharadia, S. Seth, K. Srinivasan, P. Levis, S. Katti, and P. Sinha, "Practical, real-time, full duplex wireless," in *Proc. International Conference on Mobile Computing and Networking (Mobicom 2011)*, Las Vegas, Nevada, USA, Sep. 2011, pp. 301–312.
- [Jin06] N. Jindal, "A feedback reduction technique for MIMO broadcast channels," in *Proc. IEEE International Symposium on Information Theory 2006*, Seattle, Washington, USA, Jul. 2006, pp. 2699–2703.
- [JL09] N. Jindal and A. Lozano, "De-hyping transmit diversity in modern MIMO cellular systems," in *Proc. IEEE Global Telecommunications Conference (GLOBECOM 2009)*, Honolulu, Hawaii, USA, Dec. 2009.
- [JL10] ———, "Fading models and metrics for contemporary wireless systems," in *Proc. Asilomar Conference on Signals, Systems and Computers (ASILOMAR 2010)*, Pacific Grove, California, USA, Nov. 2010, pp. 625–629.
- [KBDG12] J. Khun-Jush, P. Bender, B. Deschamps, and M. Gundlach, "Licensed shared access as complementary approach to meet spectrum demands: benefits for next generation cellular systems," in *Proc. ETSI Workshop on Reconfigurable Radio Systems*, Cannes, France, Dec. 2012.
- [KDA+10] K. Kusume, G. Dietl, T. Abe, H. Taoka, and S. Nagata, "System level performance of downlink MU-MIMO transmission for 3GPP LTE-Advanced," in *Proc. Vehicular Technology Conference (VTC-Spring 2010)*, Taipei, Taiwan, May 2010.
- [KFBV06] M. Karakayali, G. Foschini, and R. Valenzuela, "Network coordination for spectrally efficient communications in cellular systems," *IEEE Wireless Communications*, vol. 13, no. 4, pp. 56–61, 2006.
- [KH95] R. Knopp and P. Humblet, "Multiple-accessing over frequency-selective fading channels," in *Proc. Personal, Indoor and Mobile Radio Communications (PIMRC '95)*, vol. 3, Toronto, Canada, Sep. 1995, pp. 1326–1330.
- [KKLK05] H. Kim, H. Kim, J. Li, and M. Kountouris, "On the performance of limited feedback multiuser MIMO transmission in 3GPP HSDPA," in *Proc. IEEE Vehicular Technology Conference (VTC-Fall 2005)*, Dallas, TX, USA, Sep. 2005, pp. 473–476.

- [KMH+08] P. Kyösti, J. Meinilä, L. Hentilä, X. Zhao, T. Jämsä, C. Schneider, M. Narandžić, M. Milojević, A. Hong, J. Ylitalo, V.-M. Holappa, M. Alatossava, R. Bultitude, Y. de Jong, and T. Rautiainen, "WINNER II channel models; Part I channel models," IST-4-027756 WINNER II, Tech. Rep. Deliverable D1.1.2 V1.2, Feb. 2008.
- [KPL06] H. Kim, C. S. Park, and K. B. Lee, "On the performance of multiuser MIMO systems in WCDMA / HSDPA : beamforming , feedback and user diversity," *IEICE Transactions on Communications*, vol. E89-B, no. 8, pp. 2161–2169, 2006.
- [KSP+02] J. Kermoal, L. Schumacher, K. Pedersen, P. Mogensen, and F. Frederiksen, "A stochastic MIMO radio channel model with experimental validation," *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 6, pp. 1211–1226, 2002.
- [LCK+09] S.-B. Lee, S. Choudhury, A. Khoshnevis, S. Xu, and S. Lu, "Downlink MIMO with frequency-domain packet scheduling for 3GPP LTE," in *Proc. IEEE INFOCOM 2009*, 2009, pp. 1269–1277.
- [LHL+08] D. Love, R. Heath, V. K. N. Lau, D. Gesbert, B. Rao, and M. Andrews, "An overview of limited feedback in wireless communication systems," *IEEE Journal on Selected Areas in Communications*, vol. 26, no. 8, pp. 1341–1365, 2008.
- [LHS03] D. Love, R. Heath, and T. Strohmer, "Grassmannian beamforming for multiple-input multiple-output wireless systems," *IEEE Transactions on Information Theory*, vol. 49, no. 10, pp. 2735–2747, Oct. 2003.
- [LHSH04] D. Love, R. Heath, W. Santipach, and M. Honig, "What is the value of limited feedback for MIMO channels?" *IEEE Communications Magazine*, vol. 42, no. 10, pp. 54–59, 2004.
- [LJ10] A. Lozano and N. Jindal, "Transmit diversity vs. spatial multiplexing in modern MIMO systems," *IEEE Transactions on Wireless Communications*, vol. 9, no. 1, pp. 186–197, Jan. 2010.
- [LKL+12] J. Lee, Y. Kim, H. Lee, B. Ng, D. Mazzarese, J. Liu, W. Xiao, and Y. Zhou, "Coordinated multipoint transmission and reception in LTE-Advanced systems," *IEEE Communications Magazine*, vol. 50, no. 11, pp. 44–50, Nov. 2012.
- [LPT+12] E. Lähtekangas, K. Pajukoski, E. Tirola, J. Hämäläinen, and Z. Zheng, "On the performance of LTE-Advanced MIMO: how to set and reach beyond 4G targets," in *Proc. European Wireless (EW 2012)*, Poznań, Poland, 2012.
- [LR13] M. Lerch and M. Rupp, "Measurement-based evaluation of the LTE MIMO downlink at different antenna configurations," in *Proc. ITG Workshop on Smart Antennas (WSA 2013)*, Stuttgart, Germany, Mar. 2013.

- [LV89] R. Lupas and S. Verdu, "Linear multiuser detectors for synchronous code-division multiple-access channels," *IEEE Transactions on Information Theory*, vol. 35, no. 1, pp. 123–136, 1989.
- [LWKL11] S. Liu, J. Wu, C. H. Koh, and V. K. N. Lau, "A 25 Gb/s/km² urban wireless network beyond IMT-Advanced," *IEEE Communications Magazine*, vol. 49, no. 2, pp. 122–129, 2011.
- [LYC+13] C. Lim, T. Yoo, B. Clerckx, B. Lee, and B. Shim, "Recent trend of multiuser MIMO in LTE-Advanced," *IEEE Communications Magazine*, vol. 51, no. 3, pp. 127–135, Mar. 2013.
- [Mac79] V. MacDonald, "The cellular concept," *Bell System Technical Journal*, vol. 58, no. 1, pp. 15–41, 1979.
- [Mar10] T. Marzetta, "Noncooperative cellular wireless with unlimited numbers of base station antennas," *IEEE Transactions on Wireless Communications*, vol. 9, no. 11, pp. 3590–3600, Nov. 2010.
- [MAT13] MATLAB, *version 8.1 (R2013a)*. Natick, Massachusetts, USA: The MathWorks Inc., 2013.
- [MBS10] S. Max, D. Bültmann, R. Jennen, and M. Schinnenburg, "Evaluation of IMT-Advanced scenarios using the open wireless network simulator," in *Proc. International ICST Conference on Simulation Tools and Techniques (SIMUTools 2010)*, Torremolinos, Spain, Mar. 2010.
- [MCR11] C. Mehlführer, S. Caban, and M. Rupp, "Cellular system physical layer throughput: how far off are we from the Shannon bound?" *IEEE Wireless Communications*, vol. 18, no. 6, pp. 54–63, 2011.
- [MCŠ+11] C. Mehlführer, J. Colom Ikuno, M. Šimko, S. Schwarz, M. Wrulich, and M. Rupp, "The Vienna LTE simulators — enabling reproducibility in wireless communications research," *EURASIP Journal on Advances in Signal Processing*, vol. 2011, no. 1, 2011.
- [MF11] P. Marsch and G. P. Fettweis, Eds., *Coordinated Multi-Point in Mobile Communications: From Theory to Practice*. Cambridge University Press, 2011.
- [MLD11] A. Mills, D. Lister, and M. De Vos, "Understanding static intercell interference coordination mechanisms in LTE," *Journal of Communications*, vol. 6, no. 4, pp. 312–318, Jul. 2011.
- [MLKL12] S.-H. Moon, K.-J. Lee, J. Kim, and I. Lee, "Link performance estimation techniques for MIMO-OFDM systems with maximum likelihood receiver," *IEEE Transactions on Wireless Communications*, vol. 11, no. 5, pp. 1808–1816, 2012.
- [MMF97] H. Meyr, M. Moeneclaey, and S. Fechtel, *Digital Communication Receivers, Synchronization, Channel Estimation, and Signal Processing*, ser. Wiley Series in Telecommunications and Signal Processing. Wiley, 1997.

- [MN98] M. Matsumoto and T. Nishimura, "Mersenne twister: a 623-dimensionally equidistributed uniform pseudo-random number generator," *ACM Transactions on Modeling and Computer Simulation*, vol. 8, no. 1, pp. 3–30, Jan. 1998.
- [MNK+07] P. Mogensen, W. Na, I. Kovács, F. Frederiksen, A. Pokhariyal, K. Pedersen, T. Kolding, K. Hugel, and M. Kuusela, "LTE capacity compared to the Shannon bound," in *Proc. IEEE Vehicular Technology Conference (VTC-Spring 2007)*, Dublin, Ireland, Apr. 2007, pp. 1234–1238.
- [Mol06] E. Mollick, "Establishing Moore's law," *IEEE Annals of the History of Computing*, vol. 28, no. 3, pp. 62–75, 2006.
- [Mol12] A. Molisch, *Wireless Communications*, ser. Wiley-IEEE. Wiley, 2012.
- [MPR+12] P. Mogensen, K. Pajukoski, B. Raaf, E. Tirola, E. Lahetkangas, I. Z. Kovács, G. Berardinelli, L. G. U. Garcia, L. Hu, and A. F. Cattoni, "B4G local area: high level requirements and system design," in *Proc. IEEE Globecom Workshops*, Anaheim, California, USA, Dec. 2012, pp. 613–617.
- [MS96] D. Monderer and L. Shapley, "Potential games," *Games and Economic Behavior*, vol. 14, no. 1, pp. 124–143, 1996.
- [MWI+09] C. Mehlführer, M. Wrulich, J. C. Ikuno, D. Bosanska, and M. Rupp, "Simulating the long term evolution physical layer," in *Proc. European Signal Processing Conference (EUSIPCO 2009)*, Glasgow, Scotland, Aug. 2009.
- [NAV02] M. Nakamura, Y. Awad, and S. Vadgama, "Adaptive control of link adaptation for high speed downlink packet access (HSDPA) in W-CDMA," in *Proc. Wireless Personal Multimedia Communications*, Oct. 2002, pp. 382–386.
- [NC06] N. Nie and C. Comaniciu, "Adaptive channel allocation spectrum etiquette for cognitive radio networks," *Mobile Networks and Applications*, vol. 11, no. 6, pp. 779–797, 2006.
- [NNB+13] T. Nakamura, S. Nagata, A. Benjebbour, Y. Kishiyama, T. Hai, S. Xiaodong, Y. Ning, and L. Nan, "Trends in small cell enhancements in LTE-Advanced," *IEEE Communications Magazine*, vol. 51, no. 2, pp. 98–105, 2013.
- [Nok11] Nokia Siemens Networks, "2020: Beyond 4G radio evolution for the Gigabit experience," White Paper, Sep. 2011.
- [NRTV07] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, *Algorithmic Game Theory*. Cambridge University Press New York, NY, USA, 2007.

- [NST+07] M. Narandžić, C. Schneider, R. S. Thomä, T. Jämsä, P. Kyösti, and X. Zhao, "Comparison of SCM, SCME, and WINNER channel models," in *Proc. IEEE Vehicular Technology Conference (VTC-Spring 2007)*, Dublin, Ireland, Apr. 2007, pp. 413–417.
- [OR94] M. Osborne and A. Rubinstein, *A Course in Game Theory*. MIT Press, 1994.
- [PGB+11] G. Piro, L. Grieco, G. Boggia, F. Capozzi, and P. Camarda, "Simulating LTE cellular systems: an open-source framework," *IEEE Transactions on Vehicular Technology*, vol. 60, no. 2, pp. 498–513, Feb. 2011.
- [PHH+12] M. Poikselkä, H. Holma, J. Hongisto, J. Kallio, and A. Toskala, *Voice over LTE (VoLTE)*. Wiley, 2012.
- [PNG03] A. Paulraj, R. Nabar, and D. Gore, *Introduction to Space Time Wireless Communications*. Cambridge University Press, 2003.
- [PPM+07] A. Pokhariyal, K. Pedersen, G. Monghal, I. Kovács, C. Rosa, T. Kolding, and P. Mogensen, "HARQ aware frequency domain packet scheduler with different degrees of fairness for the UTRAN long term evolution," in *Proc. IEEE Vehicular Technology Conference (VTC-Spring 2007)*, Dublin, Ireland, Apr. 2007, pp. 2761–2765.
- [PSuQT13] E. Pateromichelakis, M. Shariat, A. ul Quddus, and R. Tafazolli, "On the evolution of multi-cell scheduling in 3GPP LTE / LTE-A," *IEEE Communications Surveys Tutorials*, vol. 15, no. 2, pp. 701–717, 2013.
- [PVB+09] V. Pauli, I. Viering, C. Buchner, E. Saad, G. Liebl, and A. Klein, "Efficient link-to-system level modeling for accurate simulations of MIMO-OFDM systems," in *Proc. IEEE International Conference on Communications (ICC '09)*, Dresden, Germany, Jun. 2009.
- [PWS+04] R. Pabst, B. H. Walke, D. Schultz, P. Herhold, H. Yanikomeroglu, S. Mukherjee, H. Viswanathan, M. Lott, W. Zirwas, M. Dohler, H. Aghvami, D. Falconer, and G. Fettweis, "Relay-based deployment concepts for wireless and mobile broadband radio," *IEEE Communications Magazine*, vol. 42, no. 9, pp. 80–89, 2004.
- [PWSF12] K. Pedersen, Y. Wang, B. Soret, and F. Frederiksen, "eICIC functionality and performance for LTE HetNet co-channel deployments," in *Proc. IEEE Vehicular Technology Conference (VTC-Fall 2012)*, Québec City, Canada, Sep. 2012.
- [RCMW11] M. Rupp, S. Caban, C. Mehlführer, and M. Wrulich, *Evaluation of HSDPA and LTE: From Testbed Measurements to System Level Performance*. Wiley, 2011.
- [Rin47] D. Ring, "Mobile telephony — wide area coverage," Bell Telephone Laboratories, Technical Memorandum Case 20564, Dec. 1947.

- [RSM+13] T. Rappaport, S. Sun, R. Mayzus, H. Zhao, Y. Azar, K. Wang, G. Wong, J. Schulz, M. Samimi, and F. Gutierrez, "Millimeter wave mobile communications for 5G cellular: it will work!" *IEEE Access Journal*, vol. 1, no. 1, pp. 335–349, 2013.
- [SBKN13] Y. Saito, A. Benjebbour, Y. Kishiyama, and T. Nakamura, "System-level performance evaluation of downlink non-orthogonal multiple access (NOMA)," in *Proc. IEEE Personal, Indoor and Mobile Radio Communications Conference (PIMRC 2013)*, London, Great Britain, Sep. 2013.
- [SDO+05] M. Schinnenburg, F. Debus, A. Otyakmaz, L. Berlemann, and R. Pabst, "A framework for reconfigurable functions of a multi-mode protocol layer," in *Proc. SDR Forum 2005*, Los Angeles, CA, USA, Nov. 2005.
- [SFFM99] M. Speth, S. Fechtel, G. Fock, and H. Meyr, "Optimum receiver design for wireless broad-band systems using OFDM. I," *IEEE Transactions on Communications*, vol. 47, no. 11, pp. 1668–1677, 1999.
- [SGP+13] S. Sun, Q. Gao, Y. Peng, Y. Wang, and L. Song, "Interference management through CoMP in 3GPP LTE-Advanced networks," *IEEE Wireless Communications Magazine*, vol. 20, no. 1, pp. 59–66, 2013.
- [SH72] W. B. Smith and R. R. Hocking, "Algorithm AS 53: Wishart variate generator," *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, vol. 21, no. 3, pp. 341–345, 1972.
- [Sha48] C. E. Shannon, "A mathematical theory of communication," *The Bell System Technical Journal*, vol. 27, 379–423, 623–656, Jul. 1948.
- [SHT11] C. Suh, M. Ho, and D. Tse, "Downlink interference alignment," *IEEE Transactions on Communications*, vol. 59, no. 9, pp. 2616–2626, 2011.
- [Sim07] A. Simonsson, "Frequency reuse and intercell interference co-ordination in E-UTRA," in *Proc. IEEE Vehicular Technology Conference (VTC-Spring 2007)*, Dublin, Ireland, Apr. 2007, pp. 3091–3095.
- [SPM+12] Z. Shen, A. Papasakellariou, J. Montojo, D. Gerstenberger, and F. Xu, "Overview of 3GPP LTE-Advanced carrier aggregation for 4G wireless communications," *IEEE Communications Magazine*, vol. 50, no. 2, pp. 122–130, 2012.
- [SSR11] S. Schwarz, M. Simko, and M. Rupp, "On performance bounds for MIMO OFDM based wireless communication systems," in *Proc. IEEE Workshop on Signal Processing Advances in Wireless Communications (SPAWC 2011)*, San Francisco, California, USA, Jun. 2011, pp. 311–315.

- [STB11] S. Sesia, I. Toufik, and M. Baker, *LTE - The UMTS Long Term Evolution: From theory to practice*. Wiley, 2011.
- [Str09] G. Strang, *Introduction to linear algebra*, 4th ed. Wellesley-Cambridge Press, 2009.
- [SZS07] K. Sayana, J. Zhuang, and K. Stewart, "Link performance abstraction based on mean mutual information per bit (MMIB) of the LLR channel," IEEE 802.16 Broadband Wireless Access Working Group, Tech. Rep. IEEE C802.16m-07/097, May 2007.
- [TBH08a] M. Trivellato, F. Boccardi, and H. Huang, "On transceiver design and channel quantization for downlink multiuser MIMO systems with limited feedback," *IEEE Journal on Selected Areas in Communications*, vol. 26, no. 8, pp. 1494–1504, Oct. 2008.
- [TBH08b] ———, "Zero-forcing vs unitary beamforming in multiuser MIMO systems with limited feedback," in *Proc. IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC 2008)*, Cannes, France, Sep. 2008.
- [Tel99] E. Telatar, "Capacity of multi-antenna Gaussian channels," *European Transactions on Telecommunications*, vol. 10, no. 6, pp. 585–595, 1999.
- [TG09] R. Tresch and M. Guillaud, "Cellular interference alignment with imperfect channel knowledge," in *Proc. IEEE ICC Workshops 2009*, Dresden, Germany, Jun. 2009.
- [TJC99] V. Tarokh, H. Jafarkhani, and A. Calderbank, "Space-time block codes from orthogonal designs," *IEEE Transactions on Information Theory*, vol. 45, no. 5, pp. 1456–1467, Jul. 1999.
- [Tse01] D. Tse, "Multiuser diversity in wireless networks: smart scheduling, dumb antennas and epidemic communication," in *IMA workshop on wireless networks*, Aug. 2001. [Online]. Available: <http://degas.eecs.berkeley.edu/~dtse/ima810.pdf> (visited on 2014-04-28).
- [Tse99] ———, "Forward link multiuser diversity through proportional fair scheduling," Presentation at Bell Labs, 1999.
- [Ver89] S. Verdú, "Computational complexity of optimum multiuser detection," *Algorithmica*, vol. 4, no. 1-4, pp. 303–312, 1989.
- [VKV09] P. Vandewalle, J. Kovačević, and M. Vetterli, "Reproducible research in signal processing," *IEEE Signal Processing Magazine*, vol. 26, no. 3, pp. 37–47, May 2009.
- [vRH09] C. van Rensburg and P. Hosein, "Interference coordination through network-synchronized cyclic beamforming," in *Proc. IEEE Vehicular Technology Conference (VTC-Fall 2009)*, Anchorage, USA, Sep. 2009.

- [VTL02] P. Viswanath, D. Tse, and R. Laroia, "Opportunistic beamforming using dumb antennas," *IEEE Transactions on Information Theory*, vol. 48, no. 6, pp. 1277–1294, 2002.
- [Win84] J. Winters, "Optimum combining in digital mobile radio with cochannel interference," *IEEE Journal on Selected Areas in Communications*, vol. 2, no. 4, pp. 528–539, Jul. 1984.
- [Win98] J. Winters, "Smart antennas for wireless systems," *IEEE Personal Communications*, vol. 5, no. 1, pp. 23–27, Feb. 1998.
- [WMJV96] J. Wigard, P. Mogensen, J. Johansen, and B. Vejlgaard, "Capacity of a GSM network with fractional loading and random frequency hopping," in *Proc. Personal, Indoor and Mobile Radio Communications Conference, 1996. (PIMRC 1996)*, vol. 2, Taipei, Taiwan, Oct. 1996, pp. 723–727.
- [WPS+07] N. Wei, A. Pokhariyal, T. Sørensen, T. Kolding, and P. Mogensen, "Performance of MIMO with frequency domain packet scheduling in UTRAN LTE downlink," in *Proc. IEEE Vehicular Technology Conference (VTC-Spring 2007)*, Dublin, Ireland, Apr. 2007.
- [WPS+08] N. Wei, A. Pokhariyal, T. Sørensen, T. Kolding, and P. Mogensen, "Performance of spatial division multiplexing MIMO with frequency domain packet scheduling: from theory to practice," *IEEE Journal on Selected Areas in Communications*, vol. 26, no. 6, pp. 890–900, Aug. 2008.
- [WSS06] H. Weingarten, Y. Steinberg, and S. Shamai, "The capacity region of the Gaussian multiple-input multiple-output broadcast channel," *IEEE Transactions on Information Theory*, vol. 52, no. 9, pp. 3936–3964, 2006.
- [WSS10] S. Wagner, S. Sesia, and D. Slock, "On unitary beamforming for MIMO broadcast channels," in *Proc. IEEE International Conference on Communications (ICC 2010)*, Cape Town, South Africa, May 2010.
- [XLH07] Y. Xiang, J. Luo, and C. Hartmann, "Inter-cell interference mitigation through flexible resource reuse in OFDMA based communication networks," in *Proc. European Wireless Conference (EW 2007)*, Paris, France, Apr. 2007.
- [YG06] T. Yoo and A. Goldsmith, "On the optimality of multiantenna broadcast scheduling using zero-forcing beamforming," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 3, pp. 528–541, Mar. 2006.
- [YHHM13] C. Yang, S. Han, X. Hou, and A. Molisch, "How do we design CoMP to achieve its promised potential?" *IEEE Wireless Communications*, vol. 20, no. 1, pp. 67–74, 2013.

- [YWW13] S. Ye, S. H. Wong, and C. Worrall, "Enhanced physical downlink control channel in LTE Advanced Release 11," *IEEE Communications Magazine*, vol. 51, no. 2, pp. 82–89, 2013.
- [ZM13] J. Zander and P. Mähönen, "Riding the data tsunami in the cloud: myths and challenges in future wireless access," *IEEE Communications Magazine*, vol. 51, no. 3, pp. 145–151, 2013.
- [ZPR12] H. Zhang, N. Prasad, and S. Rangarajan, "MIMO downlink scheduling in LTE systems," in *Proc. IEEE INFOCOM 2012*, Orlando, Florida, USA, Mar. 2012, pp. 2936–2940.
- [ZS07] Q. Zhao and B. Sadler, "A survey of dynamic spectrum access," *IEEE Signal Processing Magazine*, vol. 24, no. 3, pp. 79–89, 2007.
- [ZX03] Y. Zheng and C. Xiao, "Simulation models with correct statistical properties for Rayleigh fading channels," *IEEE Transactions on Communications*, vol. 51, no. 6, pp. 920–928, 2003.

Standards and Standardization Documents

- [22.803] 3GPP, "Feasibility study for Proximity Services (ProSe)," 3rd Generation Partnership Project (3GPP), TR 22.803 version 12.1.0 Release 12, Mar. 2013.
- [23.703] —, "Study on architecture enhancements to support Proximity Services (ProSe)," 3rd Generation Partnership Project (3GPP), TR 23.703 version 0.3.0 Release 12, Apr. 2013.
- [25.814] —, "Physical layer aspect for evolved Universal Terrestrial Radio Access (UTRA)," 3rd Generation Partnership Project (3GPP), TR 25.814 version 1.0.1 Release 7, Dec. 2005.
- [25.996] —, "Spatial channel model for Multiple Input Multiple Output (MIMO) simulations," 3rd Generation Partnership Project (3GPP), TR 25.996 version 6.1.0 Release 6, Sep. 2003.
- [36.101] —, "Evolved Universal Terrestrial Radio Access (E-UTRA); User Equipment (UE) radio transmission and reception," 3rd Generation Partnership Project (3GPP), TS 36.101 version 11.4.0 Release 11, Apr. 2013.
- [36.104] —, "Evolved Universal Terrestrial Radio Access (E-UTRA); Base Station (BS) radio transmission and reception," 3rd Generation Partnership Project (3GPP), TS 36.104 version 11.4.0 Release 11, Apr. 2013.
- [36.133] —, "Evolved Universal Terrestrial Radio Access (E-UTRA); Requirements for support of radio resource management," 3rd Generation Partnership Project (3GPP), TS 36.133 version 11.4.0 Release 11, Apr. 2013.
- [36.211] —, "Evolved Universal Terrestrial Radio Access (E-UTRA); Physical channels and modulation," 3rd Generation Partnership Project (3GPP), TS 36.211 version 10.0.0 Release 10, Jan. 2011.
- [36.212] —, "Evolved Universal Terrestrial Radio Access (E-UTRA); Multiplexing and channel coding," 3rd Generation Partnership Project (3GPP), TS 36.212 version 10.0.0 Release 10, Jan. 2011.
- [36.213] —, "Evolved Universal Terrestrial Radio Access (E-UTRA); Physical layer procedures," 3rd Generation Partnership Project (3GPP), TS 36.213 version 10.2.0 Release 10, Jun. 2010.

- [36.300] —, “Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Overall description; Stage 2,” 3rd Generation Partnership Project (3GPP), TS 36.300 version 10.2.0 Release 10, Jan. 2011.
- [36.306] —, “Evolved Universal Terrestrial Radio Access (E-UTRA); User Equipment (UE) radio access capabilities,” 3rd Generation Partnership Project (3GPP), TS 36.306 version 11.3.0 Release 11, Mar. 2013.
- [36.321] —, “Evolved Universal Terrestrial Radio Access (E-UTRA); Medium Access Control (MAC) protocol specification,” 3rd Generation Partnership Project (3GPP), TS 36.321 version 11.0.0 Release 11, Oct. 2012.
- [36.323] —, “Evolved Universal Terrestrial Radio Access (E-UTRA); Packet Data Convergence Protocol (PDCP) specification,” 3rd Generation Partnership Project (3GPP), TS 36.323 version 11.0.0 Release 11, Oct. 2012.
- [36.420] —, “Evolved Universal Terrestrial Radio Access Network (E-UTRAN); X2 general aspects and principles,” 3rd Generation Partnership Project (3GPP), TS 36.420 version 8.1.0 Release 8, Dec. 2008.
- [36.423] —, “Evolved Universal Terrestrial Radio Access Network (E-UTRAN); X2 Application Protocol (X2AP),” 3rd Generation Partnership Project (3GPP), TS 36.423 version 8.2.0 Release 8, Jun. 2008.
- [36.814] —, “Evolved Universal Terrestrial Radio Access (E-UTRA); Further advancements for E-UTRA physical layer aspects,” 3rd Generation Partnership Project (3GPP), TR 36.814 version 9.0.0 Release 9, Mar. 2010.
- [36.819] —, “Coordinated multi-point operation for LTE physical layer aspects,” 3rd Generation Partnership Project (3GPP), TR 36.819 version 11.1.0 Release 11, Dec. 2011.
- [36.829] —, “Enhanced performance requirement for LTE user equipment (UE),” 3rd Generation Partnership Project (3GPP), TR 36.829 version 11.1.0 Release 11, Dec. 2012.
- [36.932] —, “Scenarios and requirements for small cell enhancements for E-UTRA and E-UTRAN,” 3rd Generation Partnership Project (3GPP), TR 36.932 version 12.1.0 Release 12, Mar. 2013.
- [36.942] —, “Radio Frequency (RF) system scenarios,” 3rd Generation Partnership Project (3GPP), TR 36.942 version 11.0.0 Release 11, Oct. 2012.

- [EMD] R. Srinivasan, J. Zhuang, L. Jalloul, R. Novak, and J. Park, "IEEE 802.16m evaluation methodology document (EMD)," IEEE 802.16 Task Group m (TGm), Tech. Rep. IEEE 802.16m-08/004r5, Jan. 2009.
- [Fin09a] Finland, "Guidelines for using IMT-Advanced channel models," Tech. Rep., Jul. 2009. [Online]. Available: <http://www.itu.int/oth/R0A06000021/en> (visited on 2014-04-28).
- [Fin09b] —, "Software implementation of IMT.Eval channel model," CELTIC, Tech. Rep. Document 5D/478-E, Jul. 2009. [Online]. Available: <http://www.itu.int/oth/R0A06000022/en> (visited on 2014-04-28).
- [M.1225] ITU-R, "Guidelines for evaluation of radio transmission technologies for IMT-2000," ITU Radiocommunication Sector, Rec. ITU-R M.1225, 1997.
- [M.1645] —, "Framework and overall objectives of the future development of IMT-2000 and systems beyond IMT-2000," ITU Radiocommunication Sector, Rep. ITU-R M.1645, Jun. 2003.
- [M.2134] —, "Requirements related to technical performance for IMT-Advanced radio interface(s)," ITU Radiocommunication Sector, Rep. ITU-R M.2134, Nov. 2008.
- [M.2135] —, "Guidelines for evaluation of radio interface technologies for IMT-Advanced," ITU Radiocommunication Sector, Rep. ITU-R M.2135-1, Dec. 2009.
- [M.2135s] —, "Guidelines for evaluation of radio interface technologies for IMT-Advanced," ITU Radiocommunication Sector, Rep. ITU-R M.2135, Nov. 2008.
- [Nor03] Nortel Networks, "A discussion on angle spread calculation," 3GPP-3GPP2 SCM AHG, Tech. Rep. SCM-119-Angle Spread, Mar. 2003.
- [R1-01-0589] NEC and Telecom MODUS, "Selection of MCS levels in HSDPA," 3GPP TSG-RAN WG1, TDoc R1-01-0589, May 2001.
- [R1-030354] Samsung and Seoul National University, "Per unitary basis stream user and rate control (pu2rc)," 3GPP TSG-RAN WG1, TDoc R1-030354, Feb. 2003.
- [R1-050507] Huawei, "Soft frequency reuse scheme for UTRAN LTE," 3GPP TSG-RAN WG1, TDoc R1-050507, May 2005.
- [R1-050718] Motorola, "Simulation methodology for EUTRA UL: IFDMA and DFT-spread-OFDMA," 3GPP TSG-RAN WG1, TDoc R1-051335, Nov. 2005.
- [R1-050841] Huawei, "Further analysis of soft frequency reuse scheme," 3GPP TSG-RAN WG1, TDoc R1-050841, Sep. 2005.

- [R1-050889] Samsung, "MIMO for long term evolution," 3GPP TSG-RAN WG1, TDoc R1-050889, Aug. 2005.
- [R1-051352] —, "Simulation methodology for EUTRA uplink: SC-FDMA and OFDMA," 3GPP TSG-RAN WG1, TDoc R1-051352, Nov. 2005.
- [R1-091482] NTT DOCOMO, "Evaluation of ITU test environments," 3GPP TSG-RAN WG1, TDoc R1-091482, Mar. 2009.
- [R1-092019] Ericsson, "Summary from email discussion on calibration step 1+2," 3GPP TSG-RAN WG1, TDoc R1-092019, May 2009.
- [R1-092742] —, "Email discussion summary on calibration step 1c," 3GPP TSG-RAN WG1, TDoc R1-092742, Jun. 2009.
- [R1-110586] LG Electronics, Motorola Mobility, Intel, Marvell, Samsung, ETRI, Orange, Renesas, and Huawei, "Proposal for UE receiver assumption in CoMP simulations," 3GPP TSG-RAN WG1, TDoc R1-110586, Jan. 2011.
- [R1-111031] Nokia and Nokia Siemens Networks, "On advanced UE MMSE receiver modelling in system simulations," 3GPP TSG-RAN WG1, TDoc R1-111031, Feb. 2011.
- [R2-126087] Q. Inc., "Moving the tm5 capability," 3GPP TSG-RAN WG2, Change Request R2-126087, Dec. 2012.
- [RFC5681] M. Allman, V. Paxson, and E. Blanton, *TCP Congestion Control*, RFC 5681 (Draft Standard), Internet Engineering Task Force, Sep. 2009. [Online]. Available: <http://www.ietf.org/rfc/rfc5681.txt> (visited on 2014-04-28).
- [RP-121924] T. R. WG2, "RAN2 endorsed 36.331 CRs related to TM5," 3GPP TSG Radio Access Network, TDoc RP-121924, Dec. 2012.
- [RP-122028] Ericsson, "3GPP Work Item Description; New Carrier Type for LTE," 3GPP TSG-RAN WG1,2,4, TDoc RP-122028, Dec. 2012.
- [RP-122034] Nokia Siemens Networks, "3GPP Work Item Description; Study on 3D-channel model for Elevation Beamforming and FD-MIMO studies for LTE," 3GPP TSG-RAN WG1, TDoc RP-122034, Dec. 2012.
- [RP-130404] MediaTek, Renesas Mobile Europe, and Broadcom Corporation, "3GPP Work Item Description; Study on Network Assisted Interference Cancellation and Suppression for LTE," 3GPP TSG-RAN WG1,4, TDoc RP-130404, Mar. 2013.
- [TS36] 3GPP, *3GPP Evolved Universal Terrestrial Radio Access (E-UTRA) specification series 36.xyz*. [Online]. Available: <http://www.3gpp.org/ftp/Specs/html-info/36-series.htm> (visited on 2014-04-28).

- [TTA] TTA PG707 Evaluation Group, "C-code for report ITU-R M.2135 channel model implementation," Tech. Rep. [Online]. Available: <http://www.itu.int/oth/R0A06000024/en> (visited on 2014-04-28).
- [Wor13] Working Party 5D, "Liaison statement to external organizations—study on IMT vision for 2020 and beyond," ITU Radiocommunication Sector, Attachment 3.11 to Document 5D/300, Feb. 2013.

Cited Websites

- [boost] *Boost C++ libraries*. [Online]. Available: <http://www.boost.org/> (visited on 2014-04-28).
- [Chi] Chinese Evaluation Group, "Calibration activities in Chinese evaluation group," Chinese Evaluation Group, Tech. Rep. [Online]. Available: <http://www.itu.int/oth/R0A06000020/en> (visited on 2014-04-28).
- [Cis14] Cisco, "Cisco visual networking index (VNI): global mobile data traffic forecast update, 2013-2018," White Paper, Feb. 2014. [Online]. Available: http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white_paper_c11-520862.pdf (visited on 2014-04-28).
- [Coo] M. Cooper, *Cooper's Law*. [Online]. Available: <http://www.arraycomm.com/technology/coopers-law> (visited on 2014-04-28).
- [Fre07] Free Software Foundation, *GNU General Public License*, Jun. 2007. [Online]. Available: <http://www.gnu.org/licenses/gpl-3.0.txt> (visited on 2014-04-28).
- [Gur13] I. Gurobi Optimization, *Gurobi optimizer reference manual*, 2013. [Online]. Available: <http://www.gurobi.com> (visited on 2014-04-28).
- [IMTAphyLP] J. Ellenbeck, *IMTAphy source code hosted on launchpad.net*, Jan. 2012. [Online]. Available: <http://launchpad.net/imtaphy> (visited on 2014-04-28).
- [IMTAphyWeb] —, *IMTAphy channel model and LTE simulator documentation website*, Jan. 2012. [Online]. Available: <http://www.lkn.ei.tum.de/personen/jan/imtaphy/index.php> (visited on 2014-04-28).
- [Int10a] International Telecommunication Union, *ITU paves way for next-generation 4G mobile technologies*, Oct. 21, 2010. [Online]. Available: http://www.itu.int/net/pressoffice/press_releases/2010/40.aspx (visited on 2014-04-28).
- [Int10b] —, *ITU world radiocommunication seminar highlights future communication technologies*, Dec. 6, 2010. [Online]. Available: http://www.itu.int/net/pressoffice/press_releases/2010/48.aspx (visited on 2014-04-28).

- [Int13] —, *The world in 2013: ICT facts and figures*, Feb. 2013. [Online]. Available: <http://www.itu.int/en/ITU-D/Statistics/Documents/facts/ICTFactsFigures2013-e.pdf> (visited on 2014-04-28).
- [MKL] *Intel math kernel library*. [Online]. Available: <http://software.intel.com/en-us/articles/intel-mkl/> (visited on 2014-04-28).
- [NS3] *Ns-3 network simulator website*. [Online]. Available: <http://www.nsnam.org> (visited on 2014-04-28).
- [Oh13] Ohloh.net, *The IMTaphy open source project on Ohloh: estimated cost page*, Apr. 2013. [Online]. Available: http://www.ohloh.net/p/imtaphy/estimated_cost (visited on 2014-04-28).
- [OMP] *The OpenMP API specification for parallel programming*. [Online]. Available: <http://openmp.org/wp/> (visited on 2014-04-28).
- [Sch08] Schweizerische Eidgenossenschaft, Bundesamt für Kommunikation (BAKOM), *BAKOM Infomailing Nr. 11: Mit LTE wird UMTS schneller*, Apr. 2008. [Online]. Available: <http://www.bakom.admin.ch/dokumentation/Newsletter/01315/02251/index.html> (visited on 2014-04-28).
- [SDS+05] J. Salo, G. Del Galdo, J. Salmi, P. Kyösti, M. Milojević, D. Laselva, and C. Schneider, *MATLAB implementation of the 3GPP spatial channel model (3GPP TR 25.996)*, Jan. 2005. [Online]. Available: http://radio.aalto.fi/en/research/rf_applications_in_mobile_communications/propagation_research/matlab_scm_implementation/ (visited on 2014-04-28).
- [WIN10] WINNER+, “Calibration for IMT-Advanced evaluations,” CELTIC, Tech. Rep., May 2010. [Online]. Available: <http://projects.celtic-initiative.org/winner+/WINNER+%20and%20ITU-R%20EG%20documents/Calibration%20for%20IMT-Advanced%20Evaluations.pdf> (visited on 2014-04-28).