Part I

Motivation and Basics

1 Introduction

Patrick Marsch and Gerhard Fettweis

1.1 Motivation

Mobile communication has gained significant importance in today's society. As of 2010, the number of mobile phone subscribers has surpassed 5 billion [ABI10], and the global annual mobile revenue is soon expected to top $1 trillion [Inf10]. While these numbers appear promising for mobile operators at first sight, the major game-changer that has come up recently is the fact that the market is more and more driven by the demand for mobile data traffic [Cis10]. This is simply because Moore's law in semiconductors leads to continuously more powerful mobile devices with larger storage capacity, which in the era of Web 2.0 require regular synchronization with the Internet. Consequently, Moore's law can also be found in the increase of data rates in wireless communications, as illustrated in Fig. 1.1. The main challenge, however, is that mobile users tend to expect the fast and cheap Internet access that they are used from their fixed lines (e.g. ADSL), but anytime and anywhere while being on the move. This puts mobile operators under the pressure to respond to the increasing traffic demand and provide a more homogeneous quality of experience (QoE) over the area (often referred to as improved *fairness*), while continuously decreasing *cost per bit* - and addressing the more and more crucial issue of *energy efficiency* [FMBF10].

But how can mobile data rates and fairness be increased in general? We have to be aware that current cellular systems are mainly limited by inter-cell interference [GK00] - especially in urban areas where the rate demand is largest and hence base station deployment is dense. Here, each point-to-point communication link is characterized by a certain ratio of desired receive signal power over interference and noise power, where Shannon [Sha48] states a clear upper bound on the *capacity* of the link. This then translates to a maximum *spectral efficiency*, i.e. the maximum data rate achievable for a given bandwidth. In fact, the standard Long Term Evolution (LTE) Release 8 [McC07] uses modulation and coding schemes and link adaptation in conjunction with hybrid automatic repeat request (HARQ) that allow to approach Shannon capacity to within less than a dB at reasonable complexity [LS06]. Hence, the increasing rate demand can surely not be met by improving point-to-point links, but requires other innovations. But which further options do we have?

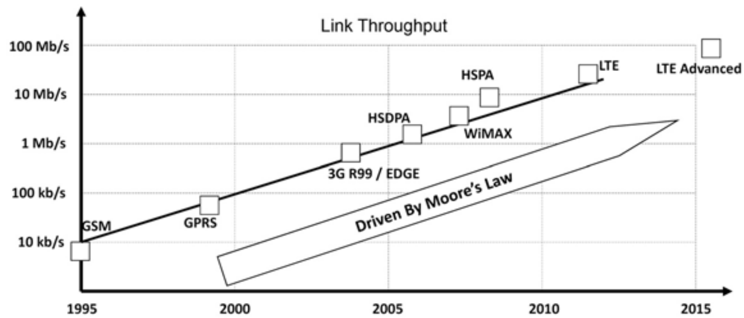

Figure 1.1 Exponential growth of data rates in mobile communications.

- **Use more spectrum**. An option which is currently already being pursued, as visible through recent auctions on spectrum becoming available via the *digital dividend*. Especially *spectrum aggregation* is of interest, i.e. the capability of radio access networks to use non-continuous blocks of spectrum. While capacity grows *linearly* with bandwidth, efficiently usable spectrum is generally a limited resource, and hence cannot be the sole source of rate growth.
- **Use more antennas** - an option that is already used since high-speed packet access (HSPA) and is a main feature of LTE. So-called multiple-input multiple-output (MIMO) techniques allow to obtain additional degrees of freedom, which can be used to spatially separate desired from interfering signals and/or for spatial multiplexing, besides yielding an additional source of diversity. In theory, capacity grows linearly in the minimum number of transmit and receive antennas [Tel99]. However, the number of antennas at base station side is usually limited due to regulatory or site rental issues, and that at the terminal side due to form factor and cost reasons. Further, practical multiplexing gains saturate at some point due to unavoidable antenna correlation.
- **Increase the degree of sectorization**. An alternative is to use *directed* base station antennas in order to obtain a larger quantity of smaller cells with less mutual interference. This is already done since global system for mobile communications (GSM), where 3-fold sectorization is typically applied, but one could principally imagine increasing sectorization [RRMF10], up to the case where each user is served with a dedicated beam.
- **Using more base stations or introducing relays and micro/femto cells** is clearly the strongest driver towards increased data rates, which also allows to improve energy and cost efficiency [MFF10]. Both relays and femto cells further allow to strongly improve indoor coverage, which is a major downside of conventional cellular systems.
- **Introduce coordination or cooperation between cells**. While most previously stated options require the deployment of new equipment, it is known from theory that interference can be overcome and even exploited if coordination or cooperation between cells is introduced. Such schemes are particularly interesting, as they require a fairly small change of infrastructure, and may

lead to a more homogeneous quality of service (QoS) distribution over the area [MKF06]. For this reason, multi-cell coordination or cooperation has been identified as a key technology of LTE-Advanced [PDF+08].

This book focuses on the latter aspect, using the term **coordinated multi-point (CoMP)**. First CoMP approaches were proposed in [BMWT00, SZ01], where the idea was to let multiple base stations jointly transmit to multiple terminals, effectively exploiting interference to obtain large gains in spectral efficiency and fairness [And05, KFVY06]. In the uplink, multiple base stations can cooperatively detect multiple terminals [WBOW00], promising similar gains [MKF06]. While the previous examples are cases of *multi-cell joint signal processing*, CoMP may also refer to schemes with a lesser extent of cooperation between base stations, for example joint scheduling or interference aware transmission and detection. In principle, it is also beneficial to let terminals cooperate [SSS+07b], but it (so far) appears difficult to explain to a mobile user why his or her handset battery is being depleted in order to enhance other users' data rates. Cooperation between terminals is hence not covered by this book.

1.2 Aim of this Book

This book provides a comprehensive overview on various CoMP techniques. It introduces information-theoretic concepts needed to understand and assess the principle degrees of freedom and gains expectable from CoMP, but also covers practical CoMP algorithms and addresses a multitude of challenges connected to their usage. A strong emphasis on implementation aspects and field trial results from the world's largest cellular research test beds gives the reader a detailed insight into the realistic potential of CoMP within the roadmap of LTE-Advanced and beyond, and the associated price and effort that have to be taken into consideration. The book provides the thorough detail required by scholars and professionals from industry or academia who aim at implementing or using CoMP themselves, but also serves as a reference book for the occasional reader.

1.3 Classes of CoMP Considered

As stated before, the term CoMP may refer to a multitude of schemes. All have in common that intra- or inter-cell interference is somehow taken into account or even exploited to enhance data rates and/or fairness. In this work, we classify CoMP schemes on one hand according to the extent of cooperation (or *information exchange*) taking place between cells:

- **Non-cooperative**, but **interference aware transceiver schemes**, where base stations or terminals adjust their transmit or receive strategy according to some knowledge on interference. This does not require explicit information exchange between cells, but the estimation of interference must be enabled through appropriate reference signal design. This class of schemes includes *single-cell multi-user* signal processing, as used in LTE Release 8 [McC07].

- **Interference coordination** schemes, where limited data is exchanged between cells for the purpose of multi-cell cooperative scheduling, multi-cell interference-aware link adaptation, or multi-cell interference-aware precoding.
- **Joint signal processing** schemes, where user data or (partially) processed transmit or receive signals are exchanged among base stations. One here considers *non-coherent* and *coherent* schemes, where the latter aim at aligning the phases of signals transmitted from or received at different antennas. As we will see, this requires precise synchronization between all involved entities.

We will later also distinguish between *decentralized* and *centralized* CoMP schemes, referring to where the subject of cooperation takes place.

1.4 Outline of this Book

The book is structured into the following 5 parts:

- **Part I - Motivation and Basics** motivates the topic from a technical and economical point of view in Chapters 1 and 2, respectively, and provides information-theoretic basics and a first insight into potential gains and trade-offs of multi-cell joint signal processing in Chapters 3 and 4.
- **Part II - Practical CoMP Schemes** introduces various specific CoMP algorithms, where Chapter 5 focusses on interference-aware transceiver schemes and interference coordination, and Chapter 6 on multi-cell joint signal processing schemes, as classified before in Section 1.3.
- **Part III - Challenges Connected to CoMP** addresses various issues regarding the usage of CoMP in practice. Chapter 7 deals with finding clusters of cells in which CoMP is performed, whereas Chapters 8 and 9 cover the crucial aspects of synchronization and channel estimation and feedback, respectively. Chapter 10 highlights practical implementation aspects such as numerical stability and scalability. Chapter 11 investigates an adaptive, situation-dependent usage of CoMP, while Chapter 12 discusses the additional backhaul infrastructure required for CoMP itself and any required signaling.
- **Part IV - Performance Assessment** discusses CoMP field trial results in Chapter 13. As field trials are usually limited to the observation of exemplary multi-point links under exemplary interference conditions, the prediction of CoMP performance in large-scale systems requires system-level simulations, for which both methodology and results are covered in Chapter 14.
- **Part V - Outlook and Conclusions** finally discusses the usage of CoMP for other purposes than rate and fairness improvements, and elaborates the usage of CoMP in conjunction with relaying or heterogeneous cellular deployment in Chapter 15. The book is then concluded in Chapter 16.

Cambridge University Press
978-1-107-00411-5 - Coordinated Multi-Point in Mobile Communications: From Theory to Practice
Edited by Patrick Marsch and Gerhard P. Fettweis
Excerpt
More information

2 An Operator's Point of View

Ralf Irmer

2.1 The Mobile Internet - A Success Story so far

When 3G was launched initially with WCDMA technology (Release 99), it was rather a disappointment with not many services being successful. Some years later, the mobile Internet took off when a number of factors came together:

- HSPA as a technological evolution of 3G with low latency and higher data rate
- Attractive flat-rate price plans by mobile operators
- Availability of mobile broadband hardware in terms of dongles and built-in 3G modules in notebooks
- Smart phones with attractive user interfaces, e.g., iPhone, Android
- Complete country-coverage with HSPA and HSPA+ by mobile operators.

This take-up of the mobile Internet generated substantial additional revenues for mobile operators, at a time when voice and text message revenues started to decline in saturated markets such as Europe. For example, Vodafone had a data revenue growth of 19% in financial year 2009/2010, with more than €4 Billion generated by non-SMS data. Today, only 11% of phones are smartphones, but by 2013 it is expected that more than a third of all active phones within the Vodafone network will be smartphones.

This data revenue growth comes along with a cost for mobile operators - namely data traffic growth. Fig. 2.1 shows the actual and projected traffic growth for Vodafone's European networks in Petabytes/year [Vod10]. It can be seen that data traffic has substantially surpassed voice traffic.

Mobile operators have some levers to cope with the growth in traffic in the short term, including:

- Technology upgrade, i.e. more efficient versions of HSPA or launch of LTE
- Cost reduction, i.e. network sharing, more efficient network operation, and exploitation of economy of scale
- Spectrum re-farming and acquisition of new spectrum
- Traffic management, i.e. enforcement of fair usage policies and launch of differentiated data bundles

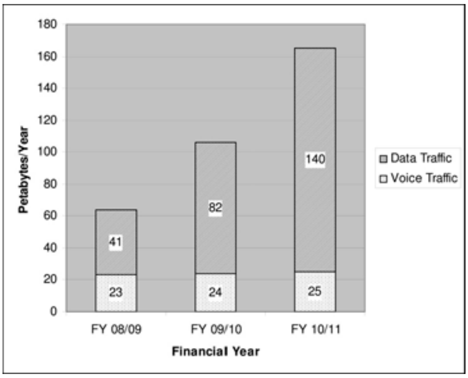

Figure 2.1 Data traffic in Vodafone's European networks in petabytes.

- Network management, i.e. building of new sites, provisioning of additional carriers or offload of traffic to femto cells or WiFi.

In the long term, however, the research community and the industry is required to come up with more fundamental approaches on how to serve mobile data at the right location in the most cost and energy efficient way.

2.2 Requirements on Future Networks and Upcoming Challenges

In 2006, a group of operators published the white paper on *Next Generation Networks beyond HSPA & EvDO* [ABG$^+$06], which lists the high-level requirements on future networks, and an accompanying document listing the detailed technical assessment criteria [IAL$^+$07]. Some of the important requirements are:

- Improved average and cell-edge spectral efficiency
- Low latency
- Simplicity, reliability and total cost of ownership
- Flat architecture

Most of the requirements are already addressed with LTE, which is being commercialized in 2010 in its first release. However, there is a need to develop LTE beyond the first release, in order to address customer and operator requirements. The challenges faced by mobile communications in the second decade of the 21st century are the following:

Exploding data volume - This is driven by attractive services, flat-rate pricing and user-friendly devices. The most prominent example is the iPhone - which resulted in a 10x traffic increase. IPTV, 3D Internet, real-time web, and cloud services will result in step changes in data consumption. IBM is predicting the generation of 16 TB/person/year by 2020. The challenge is that networks need to be structured to cope with data volume explosion without a cost or energy explosion or constant need for equipment upgrades, as illustrated in Fig. 2.2.

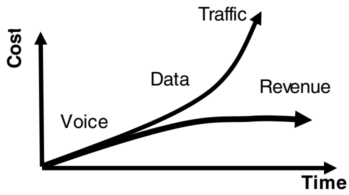

Figure 2.2 The growing gap between traffic and revenue.

Increased data rates - Driven by new services and the evolution from DSL (2 Mbps) to variants of fibre technology (100 Mbps to 1 Gbps), the user expectation of "acceptable" Internet speed will rise substantially in line with the expectation set by fibre networks and thus posing a challenge to wireless technologies.

Ubiquitous indoor coverage - Many data services are important for indoor users and people are usually within buildings. Indoor coverage is therefore important and can be either provided by copper/fibre with local radio distributions (femto cell or WiFi) or from cellular networks.

Ubiquitous outdoor coverage - For voice calls, the user expectation has moved from making calls along major roadways in the 1990s to being reachable all the time in any building. Mobile Internet based on 3G or WiFi today can only be characterized as best-effort, without "continual connectivity" whilst on-the-move and with patchy coverage in many places. In 2020, business customers and consumers will rely more on data connectivity - they will need connectivity anywhere, anytime. Coverage with a minimum guaranteed data rate and hence reliability will be a key differentiator between operators as the world moves from a "nice-to-be-connected" model to one that is "essential-to-be-connected".

There are technical innovations on the horizon to address these challenges:

- Gradual improvements of existing technologies, e.g. better MIMO modes etc.
- Active antennas, which may enable multi-element antennas
- New deployment concepts like femto cells or MetroZone networks. They require innovative backhaul solutions such as in-band and out-band backhaul or mm-wave microwave, and self-organizing principles in order to be manageable
- Miniaturized, flexible, energy-efficient base stations
- Base station cooperation concepts.

2.3 The Role of CoMP

Base station cooperation concepts (CoMP) are especially attractive since they improve the cell-edge data rate and average data rate, and are suitable to increase spectral efficiency (and hence capacity) for much more dense network deployments in urban areas and capacity hotspots. As we will see later, this increase in access capacity with CoMP concepts comes at the cost of more backhaul capacity, i.e. more communication bandwidth between base stations. However, for HSPA+ and LTE, base station sites need high-capacity backhaul (fibre or microwave) anyway, and as the cost of backhaul increases less than linearly with the backhaul capacity, this issue might not be as severe as often stated.

What are the alternatives to CoMP? Different frequency reuse, more spectrum, more sites, more antennas all are very expensive options for an operator. Thus investing into more intelligent baseband (i.e. CoMP algorithms) and backhaul with higher data rate and lower latency requirements seems to be more and more attractive when compared to the other options.

The complicated issue about CoMP concepts is that they are only partially understood from the academic perspective today, and that implementation in a standard at reasonable complexity is difficult. However, let's draw an analogy to MIMO technologies. They are commercially used today in WiFi and cellular communications, but ten years ago there was only limited understanding of MIMO, and the technology was seen by many as too complex to be commercialized.

2.4 The Role of Field Trials

Traditionally, academic innovation is evaluated using analytical models or doing statistical simulations. However, concepts such as CoMP are so complex that it is impossible to come up with models which capture all effects realistically, and to have well-calibrated simulation scenarios. Therefore, it is essential to have field trials of new technologies in an early development stage, in order to

- identify technical challenges early on
- refine simulations and analytical models
- be forced to have and end-to-end view and not pick "interesting" but non-relevant topics
- provide a proof of concept.

For CoMP technology development and evaluation, various authors of this book have set up a cluster of research test beds in Dresden and Berlin, within the EASY-C project led by Vodafone, Deutsche Telekom, Heinrich Hertz-Institute Berlin and Technische Universität Dresden. The significance of these is that enough sites are used to represent typical interference scenarios. More information on these test beds is stated in [I$^+$09], and field trial results will be presented in detail in Chapter 13.

Cambridge University Press
978-1-107-00411-5 - Coordinated Multi-Point in Mobile Communications: From Theory to Practice
Edited by Patrick Marsch and Gerhard P. Fettweis
Excerpt
More information

3 Information-Theoretic Basics

Patrick Marsch and David Tse

In this chapter, the reader is made familiar with a set of theoretical concepts to analytically capture the variety of CoMP schemes considered in this book. The reader will obtain a first understanding of the general capacity gains expectable from multi-cell joint signal processing, and the many degrees of freedom involved. The chapter introduces notation that will be reused in most parts of the book.

3.1 Observed Cellular Scenarios

Throughout the book, we generally consider (subsets of) a large cellular system as depicted in Fig. 3.1. Here, a large number of mobile terminals, or *user equipments (UEs)*, is distributed over a set of cells, where we assume that each cell is served by exactly one BS. As this is the case for most currently deployed cellular systems, we further assume that multiple BSs are grouped into so-called *sites*. Note that, differing from some other publications, we consider a *sector* to be equivalent to a *cell*. The term *cluster* is used to indicate a set of cells between which some form of CoMP may take place. Note that we assume that each UE in the system aims at transmitting or receiving dedicated information, i.e. *multi-cast* concepts are *not* covered in the book. As the number of UEs is typically significantly larger than the number of cells, UEs have to be *scheduled* to resources, i.e. to certain transmission windows. In this book, we assume that orthogonal frequency division multiple access (OFDMA) is employed as a media access technique, which allows each UE to be assigned to resources that are (under certain ideal assumptions to be discussed later) orthogonal in time and frequency. As this orthogonality allows us to simplify most of the analytical models and derivations used throughout the book, the basics of OFDMA are stated in the sequel.

3.2 Usage of OFDMA for Broadband Wireless Communications

Three fundamental challenges in mobile communications are the fact that transmission takes place over *a) a shared medium*, which is often subject to *b) rich scattering*, and to which we desire *c) simple and flexible access* of many commu-

