Master's Thesis

Pilot Contamination and Mitigation Techniques in Massive MIMO Systems

Vidit Saxena

Pilot Contamination and Mitigation Techniques in Massive MIMO Systems

Vidit Saxena
wir12vsa@student.lu.se

Department of Electrical and Information Technology
Lund University

Advisors: Dr. Eleftherios Karipidis,
Dr. Gábor Fodor

October 6, 2014

Dedicated to mamma and papa

Acknowledgement

This thesis work has been an intense learning experience for me in the methodology and rigor of wireless communications research. In a short span of six months, I have been introduced to a multitude of new ideas aimed at expanding and enriching the communication fabric of our society in the coming decades. I am grateful to Ericsson Research for giving me an opportunity to be a part of their world-leading team and working on some of the most relevant challenges in my field of study.

I offer my sincerest gratitude to my supervisors, Dr. Eleftherios Karipidis and Dr. Gábor Fodor for their guidance and support throughout the thesis. They spent an incredible amount of time and effort to discuss, suggest and review the development of this work at each stage. I thank them for making even advanced technical concepts seem understandable with their vast knowledge and encouragement. They worked hard with me to inculcate the right approach and skills for research, something that I will cherish through the rest of my career. In addition to the project work, they made sure that I was comfortable and had access to the best of resources and opportunities, for which I am exceedingly grateful.

I am grateful to all my teachers at LTH, who have given me a wealth of knowledge and kindled my interest in wireless communication research. I am especially thankful to Prof. Ove Edfors, who first introduced me to massive MIMO systems, and has been a constant source of inspiration with his outstanding work and ideas in this field.

My parents have been a constant source of support and joy while writing this thesis, and especially my mother, forever ready to solve any technical difficulties that I might face. I am grateful for their support and to my brother and his wife for bringing so much cheer to the family. Finally, I am thankful to the friends at the university and at home, who have filled the weekends with fun and laughter, making my stay in Stockholm a charming experience.

Abstract

A multi-antenna base station (BS) can spatially multiplex a few terminals over the same bandwidth, a technique known as multi-user, multiple-input multiple-output (MU-MIMO). A new idea in cellular MU-MIMO is the use of a large excess of BS antennas to serve several single-antenna terminals simultaneously. This so-called "massive MIMO" promises attractive gains in spectral efficiency with time-division duplex operation. Within a cell, the BS estimates the channel from mutually orthogonal reverse-link pilot sequences to formulate a receiver for the reverse link and (assuming reciprocity) a precoder for the forward link. The channel coherence is typically constrained in time as well as frequency, leading to a trade-off between the resources spent on pilots and those available for data symbols. This pilot overhead can be reduced by reusing pilot sequences in nearby cells, however this potentially introduces interference in the channel estimation phase, the so-called "pilot contamination" effect.

In this thesis, we study the impact of pilot contamination in realistic environments and investigate schemes to mitigate it. We evaluate the mean squared error (MSE) of channel estimates in case of a plain-vanilla least-squares (LS) estimator and a minimum MSE (MMSE) estimator that exploits prior knowledge of second-order channel statistics. Next, we introduce a pilot open-loop power control (pilot OLPC) scheme to improve the SINR-fairness of received pilot signals at the BS. We evaluate the effect of relaxing the pilot reuse factor and also implement a soft pilot reuse (SPR) scheme to distribute pilot sequences efficiently. To study the trade-off between pilot and data symbols, we evaluate the achievable rate in forward link with maximum-ratio and zero-forcing precoding at the BS. We evaluate an inter-cell coordination scheme that exploits prior knowledge of all cross-channel covariance matrices to reuse pilots among spatially well-separated terminals.

We simulate a 21-cell MU-MIMO setup with up to 100-antenna BSs and up to 24 single-antenna terminals per cell in an outdoor urban macro environment. We find that pilot reuse 1 causes severe impairment of the channel estimates, which can be improved with pilot OLPC. Pilot reuse 1/3 effectively mitigates pilot contamination, and can improve the achievable rate in the forward link. SPR also mitigates contamination but with a smaller increase in pilot overhead. Inter-cell coordinated pilot allocation, implemented using a greedy approach, provides gains over random allocation only for the initial few pilots. In general, maximum ratio precoding is more robust against pilot contamination than zero-forcing.

Table of Contents

Introduction

Wireless technology has become the primary enabler of mobility and ubiquitous network access over the past decade. The demand for higher peak data rates and anytime, anywhere connectivity has been driving the rapid developments in cellular technology. There is a need to serve a host of data-intensive applications dominated by video streaming, real time services like navigation, and graphics-heavy social media interfaces on hand-held devices. Over the past decade multi-input, multiple-output (MIMO) has emerged as one of the prime technologies for achieving high data rates and spectral efficiency in wireless communication systems. The latest wireless communication standards rely on MIMO air interface to achieve cost-effective deployment of gigabit links in wireless local area networks as well as commercial cellular networks. However, the gains with MIMO are achieved at the cost of increased processing complexity and hardware costs. The rapid developments in hardware technology over the past decade have acted as enablers for MIMO by facilitating large-scale production of low-cost chips with small form factors. The current state of the art in cellular technology (LTE Release 10) allows for 8 antenna ports at the base station (BS) and an equal number of antenna ports at the terminal. Research efforts to develop a technology that increases spectral efficiency while the bulk of processing complexity is handled at the BS has led to the evolution of massive MIMO. This solution aims to scale the BS array size to hundreds of antennas and spatially multiplex a few tens of terminals, which can increase cell throughput by an order of magnitude while improving energy efficiency and simplifying radio access.

1.1 Background

Two major challenges to communication over a wireless channel of given bandwidth are: *fading*, the variation of signal power over time and *interference* caused by other devices transmitting over the same bandwidth at the same time. These factors affect the signal to interference-and-noise ratio (SINR) at the receiver and place an ultimate limit on the achievable rate ($\mathcal{R}$) with an arbitrarily low error rate. In case of Gaussian distribution of the thermal noise and interference power, $\mathcal{R}$ has been shown to increase linearly with the bandwidth and logarithmically with the SINR.

The rate of information transmitted per unit bandwidth is known as the *spec-*

tral efficiency of a link. Since spectrum is a limited and expensive resource, spectral efficiency is one of the most important metrics to keep in mind while deploying a cellular system (another crucial metric being energy efficiency which we discuss briefly later in the thesis). Since spectral efficiency only grows logarithmically with the SINR, a large increase in signal power might result in only moderate gains in the spectral efficiency of the link. Furthermore as signal power increases, the interference caused on other links increases proportionally, limiting the improvement of spectral efficiency of the overall system. Clearly, we need alternate means to improve the spectral efficiency, either by reducing the interference or by creating additional orthogonal channels within the given spectrum.

The key approach used by MIMO to improve spectral efficiency is by (i) processing signals coherently at multiple transmitter/ receiver ports to improve the received signal power (*array gain, diversity*), (ii) interference cancellation, or (iii) transmitting independent data streams over spatially separate links (*spatial multiplexing*) [1]. These approaches are complementary, and the optimal approach in terms of maximizing $\mathcal{R}$ depends on the several factors. In a cellular scenario where the BS serves a multitude of UEs over the same time-frequency resources, spatially separated streams can be used to transmit/receive data to/from each terminal. This technique is termed multi-user MIMO (MU-MIMO), also known as space division multiple access (SDMA). With excess degrees of freedom at the BS array, spectral efficiency of MU-MIMO can increase linearly with the number of terminals served under the usual channel conditions.

To achieve gains in $\mathcal{R}$ with MU-MIMO, the BS requires fairly accurate knowledge of the forward- and reverse-link channels for each terminal. A common approach for estimating a channel is by transmitting a known sequence of symbols (a "pilot sequence") and noting the effect of the channel on this sequence at the receiver. The channel estimate is useful only over limited time and frequency interval over which the channel can be assumed constant, known as its coherence interval. Obtaining channel state information (CSI) constitutes an overhead, since the resources spent on obtaining CSI limits the amount of resources available for receiving or transmitting data within the coherence interval. This trade-off between pilot and data resources is an important research problem to determine the optimal parameters for maximizing the cell throughput.

The addition of more antennas at the BS enables it to focus the signal into smaller regions of space at the location of desired terminal, thus improving the received signal power and mitigating interference. This motivates the investigation of asymptotically infinite number of antennas at the BS in a multi-cell system, and evaluation of its impact on the $\mathcal{R}$ [2]. It is found that in the case of infinitely many antennas, the system throughput and energy efficiency increases dramatically. The effects of noise and interference vanish with simple linear processing, and theoretically the performance is only limited by the contamination of channel estimate as a result of reusing the pilots across cells.

1.2   Thesis Contributions

The problem of pilot contamination has received widespread attention within the ongoing study of massive MIMO [3][4]. Theoretical analysis has identified pilot contamination as the fundamental limit on the throughput of massive MIMO systems [2]. With this thesis, our aim is to evaluate the scenarios under which pilot contamination becomes a significant problem in realistic systems, and study techniques that mitigate its impact. We simulate a multi-cell network in an urban layout, where the elevated BS in each cell serves a number of outdoor terminals moving at vehicular speeds.

First, we investigate the worst-case pilot contamination scenario, when the same set of mutually orthogonal pilot resources are reused in each cell. We evaluate a least-squares (LS) channel estimator that allows for simple implementation and does not require any prior knowledge of channel statistics. We observe that this approach leads to significant contamination of channel estimates. Next, we investigate a minimum mean square error (MMSE) channel estimator assuming perfect knowledge of channel covariance matrices. We find that the MMSE estimator improves the channel estimates, albeit with significantly higher complexity.

Next, we implement a pilot open loop power control (pilot OLPC) scheme to improve the SINR-fairness of received pilot signals at the BS. We observe that this scheme improves LS channel estimates for several terminals, while also reducing the total transmit energy. A simple technique to further mitigate pilot contamination is to use orthogonal pilot sequences in nearby cells. We find this scheme provides nearly contamination-free channel estimates, at the cost of higher pilot overhead. The trade-off between using more resources for pilot-based channel training versus data is evaluated in terms of per-cell sum rate. To reuse pilot resources more efficiently, we develop a soft pilot reuse scheme and note the impact on the rate.

Finally, the performance of MMSE channel estimator can be improved by reusing pilots among spatially well-separated terminals. This can be accomplished by exploiting the information in channel covariance matrices and slow-rate inter-cell coordination. We investigate this technique in context of a greedy allocation algorithm, and observe that it surprisingly does not provide overall rate gain when compared with random pilot allocation.

1.3   Thesis Organization

In Ch. 2, we describe the cellular model in the context of this thesis and present an overview of the wireless channel and multiple antenna techniques. Also, we describe MU-MIMO operation and summarize the theoretical underpinnings of massive MIMO. In Ch. 3, we discusses pilot-based channel estimation techniques and describe the issue of pilot contamination in multi-cell systems. We investigate the pilot-data trade-off and arrive at an expression for the achievable rate per terminal and the per-cell sum rate. In Ch. 4, we describe methods to combat pilot contamination and the conditions under which each technique is expected to give most gains. Ch. 5 presents the simulation setup and numerical results of proposed methods. The thesis concludes with key takeaways and discussion of potential future work in Ch. 6.

Cellular Communication with Multi-user MIMO

In Sec. 2.1, we introduce the wireless propagation environment in context of an urban macro setup (elevated BS in an urban area serving terminals located upto several hundred meters away) that is studied in the rest of the thesis. We observe that with multiple scatterers and mobile terminals, the channel may be assumed approximately constant only over a limited time and bandwidth. We discuss the standard channel models that are used for simulation studies, and discuss the channel correlation at elements of an antenna array. Next, in Sec. 2.2 We summarize multiple antenna techniques to improve spectral efficiency, and explain the MU-MIMO setup. Moreover, we present expressions for linear, forward- and reverse-link processing at the BS. Finally in Sec. 2.3 we introduce the concept of massive MIMO as an natural extension of the MU-MIMO setup. We conclude this chapter with an outline of the challenges with massive MIMO and the practical scenarios where it might be feasible to implement such a system.

2.1   The Wireless Channel

Radio signals propagate through space in the form of electromagnetic waves. The radio wavefront spreads out with distance, attenuating the power density of the signal. In addition, outdoor channels contain several interfering objects (*scatterers*) such as trees, buildings and vehicles that cause the signal to be absorbed, reflected, or diffracted (Fig. 2.1). The interaction of scatterers with radio signal leads to multiple propagation paths between the transmitter and the receiver, with each path having a different propagation delay and attenuation.

In free space, the attenuation of signal power is solely on account of expansion of the signal wavefront and is known as *free-space pathloss*. In the far field of the transmitting antenna, the free space pathloss is described by an *inverse square law*, which implies that the power density of the signal decreases with the square of distance from transmitting antenna (pathloss exponent of 2). During transmission over a ground plane, as is normally the case with cellular systems, part of the radio signal gets reflected by the earth's surface and interferes with the primary wavefront. For such a scenario, the power density decreases with a pathloss exponent of 4. In practical urban environments, the pathloss exponent is

found to be between 2.5 and 6, depending on the topography and environmental factors such as foliage, moisture content etc. [5]

Large scatterers like buildings, hillocks etc severely attenuate the radio signal, creating shadow zones that span an area much larger than the wavelength of the signal carrier. This phenomenon of *shadow fading* causes *large-scale* attenuation for the duration that the terminal remains within the shadow of the scatterer. However, the wavefront undergoes diffraction at the edges of the scatterer, and some fraction of signal power is available to the terminal even when the primary wavefront might be completely obstructed. In a typical urban setup, the are several large structures that might cause this kind of fading during communication with a mobile terminal. Practical measurements have found that the logarithm of attenuation caused by shadow fading from several scatterers closely follows a normal distribution, so it is also known as *lognormal* fading.

Figure 2.1: Radio Propagation

Scattering causes the signal to arrive at the receiving antenna via several paths, with each path having a unique attenuation and phase shift. These *multipath components* (MPCs) add up constructively or destructively at the receiver, depending upon their path length. The fading due to superposition of multipath components varies with distances of the order of a few wavelengths and is therefore called *small-scale fading*. If there is a large number of uncorrelated scatterers and no line of sight, as is usually the case in heavily built-up urban areas, the in-phase and quadrature components of the received signal with small-scale fading can be assumed to follow independent zero-mean Gaussian distributions, also called Rayleigh fading.

The received signal power varies significantly with different spatial locations on account of fading. In an urban setup, the scatterers or the receiver might be in motion with respect to the transmitter, causing variations in the received signal with time. This movement causes Doppler shift in the frequency of MPCs, which is discussed in the next section. The combined effect of pathloss, large- and small-scale fading, and Doppler shift is the primary challenge in communication over a wireless channel. Several techniques have been devised over the past decades, and

intense research is ongoing for improving the throughput and reliability. We study a promising technique that, under certain assumptions and channel conditions, promises attractive gains over existing wireless systems.

Channel Coherence

The small-scale fading introduced above is caused by superposition of MPCs which interfere constructively or destructively depending upon their phase at the receiver. The phase shift of the i^{th} multipath component from initial phase ψ_0 is directly proportional to its path length Δ_i and the frequency of operation f_c,

$$\psi_i - \psi_0 = \frac{2\pi\Delta_i f_c}{c}, \tag{2.1}$$

where c is the speed of light in air.

It follows that in a multipath channel, the instantaneous phase shift of MPCs and hence the small scale fading varies with frequency. The bandwidth over which the channel can be assumed approximately constant is called its coherence bandwidth B_c. If B_c is less than the signal bandwidth, the channel is said to be *frequency selective*, otherwise it is *frequency flat*. The coherence bandwidth is inversely proportional to root mean square (RMS) of path lengths, Δ, and can be approximated as

$$B_c \approx \frac{c}{\Delta}. \tag{2.2}$$

If a scatterer or the receiver is moving relative to the transmitter, the frequency of the corresponding MPC undergoes a Doppler shift. The magnitude of Doppler shift depends upon the carrier frequency and the velocity of the receiver relative to the incoming wavefront. Superposition of components with independent Doppler shifts results in broadening of the spectrum of the received signal. This phenomenon, known as *frequency dispersion*, causes the channel response to vary with time (*time-selective* fading). The channel can be assumed approximately constant only over a coherence time T_c that is inversely proportional to the velocity v and can be approximated as

$$T_c \approx \frac{c}{4vf_c}. \tag{2.3}$$

In outdoor urban environments the channel between a BS and a mobile terminal suffers from both frequency and time selectivity. This places limits on the duration and bandwidth over which the channel can be assumed approximately constant, known as its *coherence interval* (Fig. 2.2). The channel estimates obtained via training are valid only within this interval, and can be the exploited by the multi-antenna BS for improving the SINR at the receiver. For communication over the Long-Term Evolution (LTE) forward-link grid, the coherence interval can be measured in terms of coherent time-frequency elements. A single element in this grid consists of a symbol of duration 71.4 μs and a bandwidth of 15 kHz. The number of coherent time- and frequency-resources τ_c is then obtained as

$$\tau_c = \left(\frac{B_c}{15kHz}\right)\left(\frac{T_c}{71.4\mu s}\right). \tag{2.4}$$

Figure 2.2: Coherence interval over a LTE forward-link grid

We approximate the value of τ_c for some usual and extreme scenarios in Table 2.1. In case of an urban macro environment, typical distance between the BS and the cell border is of the order of several hundred meters. We Assume that the worst-case difference in path length of MPCs is due to reflections at the cell border, causing frequency selectivity as discussed above. We assume a Δ of 1667 m for our analysis, which gives an integral value of $B_c = 180$ kHz at an operating frequency of 2 GHz. We evaluate τ_c at several terminal speeds, ranging from pedestrian (5 kmph) to vehicular (50 kmph), and extreme speeds (250 kmph). We observe that at 2 GHz, τ_c is expected to be approximately 500 and 250 resource elements for medium and high vehicular speeds respectively. This is relevant from an urban macro scenario, to set conservative limits on the assumed channel coherence interval.

In the LTE forward-link grid, the symbols for each terminal are stacked onto orthogonal subcarriers using a $N = 12$-point IFFT operation, generating an OFDM symbol that spans N subcarriers and in frequency 71.4 μs in time. A set of 7 consecutive OFDM symbols is known as a resource block (RB). Communication with each terminal is scheduled by allocating them one or more RBs in a given time interval [6]. An RB pair therefore comprises 14 OFDM symbols spanning $14 \times 12 = 168$ time-frequency elements, which is that largest number of RBs for which channel can be assumed coherent at 2 GHz and vehicular speeds. In the rest of this thesis, we assume forward- and reverse-link transmission over a channel coherence interval of a LTE RB pair.

	$f_c = 1$GHz	$f_c = 2$GHz	$f_c = 6$GHz	$f_c = 15$GHz
pedestrian	$\tau_c = 10000$	$\tau_c = 5000$	$\tau_c = 1600$	$\tau_c = 600$
10 km/h	$\tau_c = 5000$	$\tau_c = 2500$	$\tau_c = 800$	$\tau_c = 300$
50 km/h	$\tau_c = 1000$	$\tau_c = 500$	$\tau_c = 160$	$\tau_c = 60$
100 km/h	$\tau_c = 500$	$\tau_c = 250$	$\tau_c = 80$	$\tau_c = 30$
250 km/h	$\tau_c = 200$	$\tau_c = 100$	$\tau_c = 32$	$\tau_c = 12$

Table 2.1: Impact of carrier frequency and terminal speed on channel coherence

Channel Models

Actual wireless channels are complex and difficult to represent accurately. For simulation studies, empirical models have been developed based on extensive measurements that approximate the most common communication scenarios. In general, the channel coefficient between a transmit and receive antenna is modeled by a complex random variable that models pathloss, shadowing and small-scale fading effects described previously. The instantaneous magnitude and phase of the channel coefficient represents the amplitude and phase of channel's frequency response respectively.

For simulation studies, spatial channel models (SCMs) have been developed by the 3rd generation partnership project (3GPP) and international telecommunication union (ITU) that model various urban and rural propagation scenarios. The ITU-R IMT-Advanced channel model is a stochastic model based on the scenario geometry [7]. The model includes information about angle of arrival (AoA) as well as angle of departure (AoD), the so-called double-directional channel model [8]. It specifies the directions, amplitudes and phases for several *rays* (plane-waves) instead of spatial location of the scatterers. The instantaneous parameters are determined stochastically based on statistical distributions extracted from actual channel measurements for several well-known scenarios. The location, geometry and pattern of antennas can be decided according to specific scenario of the simulation study. These rays are superimposed at the location of antennas in the simulation setup, and the effects of delay, power, and angular parameters are evaluated to obtain the channel coefficients at several instants in time. Moreover, the superposition of rays produces the effects of correlation between antenna elements, temporal fading and Doppler spectrum at the transmitter as well as at the receiver.

The urban macro model (UMa) targets coverage for pedestrian and vehicular users, with non-line of sight (NLOS) as the dominant mode of propagation. The predominant scatterers are buildings, which are usually assumed to be placed in a Manhattan-grid layout. The mobile terminal is assumed to be located outdoors at ground level, while the BS is elevated to a height greater than the buildings in the vicinity. The detailed expressions for pathloss are provided by the ITU model in [7], and are used to simulate the scenario along with other parameters in Ch. 5.

We consider a setup with L time-synchronized cells containing K single-antenna terminals each. The BS in each cell is equipped with an M-antenna array that communicates with the terminal using N-subcarrier OFDM over the same time-frequency interval. The mathematical analysis is independent of the BS array geometry. Later, we simulate a uniform linear array (ULA) as a suitable candidate for deployment in an urban macro scenario. The complex propagation coefficient for n^{th} subcarrier between m^{th} BS antenna in the j^{th} cell and k^{th} terminal in l^{th} cell can denoted as g_{nmjkl}, where $n = \{1, 2, ..., N\}$, $m = \{1, 2, ..., M\}$, $j = \{1, 2, ..., L\}$, $k = \{1, 2, ..., K\}$, and $l = \{1, 2, ..., L\}$.

Spatial Correlation

The complex channel coefficient between m^{th} BS antenna in the j^{th} cell and k^{th} single-antenna terminal within the same cell can be simplified to g_{mjk}, where we

drop the dependence on subcarrier index n for simplicity. We denote the $M \times K$ channel matrix between the BS and terminals of the j^{th} cell by $\mathbf{G}_j$,

$$
\begin{aligned}
\mathbf{G}_j &= [\mathbf{g}_{j1}, \mathbf{g}_{j2}, ..., \mathbf{g}_{jK}], \\
\text{where } \mathbf{g}_{jk} &= [g_{1jk}, g_{2jk}, ..., g_{Mjk}]^T,
\end{aligned} \tag{2.5}
$$

and $\mathbf{g}_{jk}$ is the $M \times 1$ channel vector between the BS and the k^{th} terminal within that cell. The channel characteristics described in Sec. 2.1 present Rayleigh distribution as a model for small scale fading. If the conditions for this model are valid, g_{mjk} can be modeled as a zero mean, circularly symmetric complex Gaussian (ZMCSCG) random variable. Additionally, if the signal at each antenna is uncorrelated, the elements of $\mathbf{g}_{jk}$ are independent and identically distributed (i.i.d.) and the channel is said to be *spatially white*.

$$
\begin{aligned}
E\{g_{mjk}g_{njk}\} &= 0 \ if \ m \neq n, \\
\mathbf{g}_{jk} &= \mathbf{g}_w,
\end{aligned} \tag{2.6}
$$

where $\mathbf{g}_w$ is the spatially white channel.

The correlation of the signal across an antenna array depends upon the angular spread of its MPCs and the antenna spacing. If the angular spread is large, a smaller antenna spacing is sufficient to ensure uncorrelated signal at each antenna element. A mobile terminal in a typical urban setup receives signal from scatterers located in several directions. Therefore, antenna spacing as low as 0.5λ might be enough to ensure uncorrelated signals, but the number of antennas is limited by the terminal size. For an elevated BS, the required spacing is usually of the order of several tens of wavelengths due to narrow angular spread [9]. The spatial correlation between elements of $\mathbf{g}_{jk}$ can modeled using its $M \times M$ covariance matrix $\mathbf{R}_{jk}$,

$$
\begin{aligned}
\mathbf{g}_{jk} &= \mathbf{R}_{jk}^{1/2}\mathbf{g}_w, \\
\mathbf{R}_{jk} &= E\{\mathbf{g}_{jk}\mathbf{g}_{jk}^H\},
\end{aligned} \tag{2.7}
$$

where ^{H} represents conjugate and transpose (Hermitian) operator. The off-diagonal elements of $\mathbf{R}_{jk}$ model the correlation between the antenna elements, while the diagonal elements model pathloss. The covariance matrix varies slowly as compared to small scale fading and can be approximated from the channel estimates. The knowledge of spatial channel correlation can be exploited for improvement in several areas such as channel estimation [10], cell selection [11] and transmit/receive beamforming [12].

2.2   Multi-user MIMO

In the first three generations of cellular technology, the BS served multiple terminals by separating them in time, frequency or code. Each terminal was assigned a unique fraction of spectrum resources for communication over the forward- and

reverse- links, to minimize intra-cell interference. A multi-antenna BS opens up the *spatial* dimension that allows it to discriminate the signals to/from each terminal based on its location, known as MU-MIMO. This enables each terminal to use all available spectrum resources, improving the throughput without the need for additional (expensive) resources. The hardware cost involved with MU-MIMO is the need to place additional BS antennas at the locations where we wish to transmit/receive the signal. Thus, the available spatial *degrees of freedom* at the BS are limited by the number of antennas.

A multi-antenna transmitter can *precode* the signal with a complex weight vector such that the radiated energy from each antenna adds constructively or destructively in desired directions. This approach, called *transmit beamforming*, can be used to maximize the signal power at the receiver or place nulls in the direction of interferers. The optimal beamforming weights depend on the instantaneous amplitude and phase of the channel. Analogously, a multi-antenna receiver may exploit channel knowledge for *receive beamforming* to maximize signal power and minimize the interference power.

The MU-MIMO setup of interest consists of a BS with M antennas serving K single-antenna terminals ($K \leq M$), over the same time-frequency resources. The BS exploits channel knowledge for transmit and receive beamforming to create spatially separate data stream for each terminal. The data streams function as independent SISO links under favorable channel conditions, and can linearly increase the spectral efficiency with the number of terminals served. However, the benefits of this *spatial multiplexing* in terms of spectral efficiency critically depends on the array size and the accuracy of channel estimates at the BS. We consider a multi-cell time-synchronized network where the BS in each cell spatially multiplexes several single-antenna terminals. We assume that the j^{th} BS has access to channel estimates $\widehat{\mathbf{g}}_{jk}$ for all the k terminals it serves, $k \in \{1, 2, ..., K\}$. The techniques to obtain these estimates are the subject of Ch. 3. We consider only linear precoders/receivers within the context of this thesis.

(a) SISO

(b) MU-MIMO
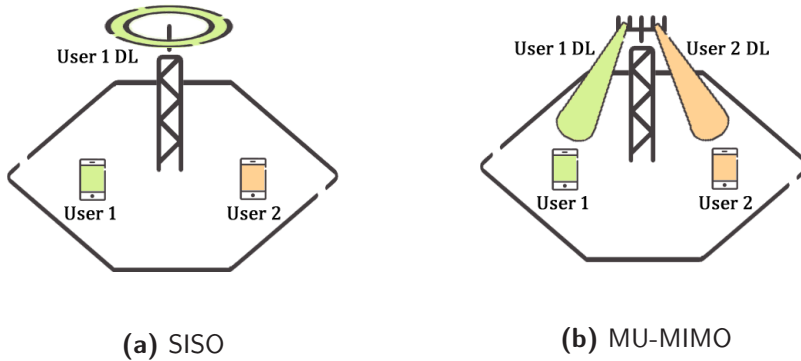
Figure 2.3: Forward-link data transmission

Forward-Link Data Transmission

In the forward link, the BS broadcasts the signal to each terminal over a shared channel, leading to intra-cell interference. The degrees of freedom at the BS can be exploited to increase the received signal power at each terminal, or reduce the interference. The forward link of the MU-MIMO cell is characterized by an $M \times K$ channel matrix $\mathbf{G}_j$ introduced in Sec. 2.1. The BS constructs its (noisy) version of the channel matrix, $\widehat{\mathbf{G}}_j = [\widehat{\mathbf{g}}_{j1} \ \widehat{\mathbf{g}}_{j2} \ ... \ \widehat{\mathbf{g}}_{jK}]$ using the channel estimates. The data symbols to be transmitted to each terminal are denoted by a $K \times 1$ vector $\mathbf{a}_j$. The BS precodes the data vector with a weighing matrix $\mathbf{W}_j$ as shown in Fig. 2.4. The signal transmitted from the BS can then be expressed as

$$
\mathbf{y}_j = \sqrt{\frac{P_{\mathrm{BS}}}{MK}} \ \mathbf{W}_j \begin{bmatrix} a_{j1} \\ a_{j2} \\ \vdots \\ a_{jK} \end{bmatrix}, \tag{2.8}
$$

$$
\text{where } \mathbf{W}_j = [\mathbf{w}_{j1} \ \mathbf{w}_{j2} \ ... \ \mathbf{w}_{jK}],
$$

$$
\sum_{k=1}^{K} \|\mathbf{w}_{jk}\| = 1,
$$

and P_{BS} is the total BS transmit power over the channel bandwidth. The data symbols are assumed to be unit-power and $\| \ \|$ denotes the Euclidean norm of a vector, and the Frobenius norm of a matrix.

The weight vectors are designed to satisfy some requirements on the received signal power or interference at each terminal. Maximum ratio transmission (MRT, also known as conjugate beamforming) is an approach that focuses transmitted signal power to the location of desired terminal, while not considering its impact on interference at rest of the terminals. MRT is therefore an SNR-maximizing technique. On the other hand, zero-forcing (ZF) design places nulls at the location of rest of the terminals to minimize interference at those terminals. However, it transmits the signal in all other directions, a fraction of which may be received by the desired terminal. In general, MRT is preferable for noise-dominated regime, while ZF is preferred for high-SNR regime where interference dominates the SINR.

The weight vectors for MRT transmission are given by Eq. 2.9, where * denotes complex conjugation. This design matches the k^{th} weight vector to the channel between the BS and the corresponding terminal, thus maximizing the received signal power. A fraction of the transmitted signal power also appears at the l^{th} ($\neq k$) terminal, depending on the inner product between the k^{th} weight vector and the channel for the l^{th} terminal. The power allocated to each link is normalized such that the total BS transmit power remains unchanged. The terminals close to the BS have significantly stronger links and are therefore allocated a higher

fraction of the total transmit power, compared to cell-edge terminals.

$$\mathbf{w}_{jk}^{MRT} \quad = \quad \frac{\widehat{\mathbf{g}}_{jk}^{*}}{\|\widehat{\mathbf{G}}_{j}\|}, \tag{2.9}$$

$$\mathbf{w}_{jk}^{ZF} \quad = \quad \frac{\widehat{\mathbf{g}}_{jk}^{\dagger}}{\|\widehat{\mathbf{G}}_{j}^{\dagger}\|}. \tag{2.10}$$

The ZF weight vectors are described by Eq. 2.10 where $\widehat{\mathbf{g}}_{jk}^{\dagger}$ is the k^{th} column of the pseudo-inverse matrix $\widehat{\mathbf{G}}_{j}^{\dagger}$ given by $\widehat{\mathbf{G}}_{j}^{\dagger} = (\widehat{\mathbf{G}}_{j}^{H}\widehat{\mathbf{G}}_{j})^{-1}\widehat{\mathbf{G}}_{j}^{H}$, and $^{-1}$ denotes matrix inversion. With this design, the k^{th} weight vector steers nulls to the location of each l^{th} ($\neq k$) terminal to minimize interference, while transmitting everywhere else. If the k^{th} terminal is located close to any l^{th} terminal, the received power is small compared to spatially separate terminal placement.

Figure 2.4: Forward-link data transmission with precoding in MU-MIMO

The received signal at a terminal suffers from intra- and inter-cell interference. As discussed above, the magnitude of intra-cell interference depends on the precoding technique. Also, a larger M implies a more number of elements in the weight vector and consequently, more effective precoding. We stack the data symbols received by the terminals of j^{th} cell in a $K \times 1$ vector $\mathbf{x}_j$,

$$\mathbf{x}_j \quad = \quad \mathbf{G}_j^T\mathbf{y}_j + \sum_{l=1,l\neq j}^{L} \mathbf{G}_{jl}^T\mathbf{y}_l + \mathbf{n}_j, \tag{2.11}$$

where ^{T} is the transpose operator, $\mathbf{G}_{jl}$ is the $M \times K$ channel between terminals of the j^{th} cell and the l^{th} BS, and $\mathbf{n}_j$ is the thermal noise at each terminal. The first term in Eq. 2.11 includes the desired signal as well as the intra-cell interference and the second term accounts for inter-cell interference from the $L - 1$ other cells.

In Eq. 2.12, we present the forward-link SINR $\Gamma_{jk}^{\mathrm{FL}}$ at k^{th} terminal of the j^{th} cell. The numerator is proportional to the received signal power. The first and second terms of the denominator are proportional to intra- and inter-cell

interference respectively. The last term corresponds to for thermal noise at the terminal, and accounts for the proportionality constant $\frac{P_{\mathrm{BS}}}{MK}$ for all terms.

$$\Gamma_{jk}^{\mathrm{FL}} = \frac{|\mathbf{g}_{jk}^T \mathbf{w}_{jk}|^2}{\sum_{i \neq k} |\mathbf{g}_{jk}^T \mathbf{w}_{ji}|^2 + \sum_{l \neq j} \sum_{i=1}^{K} |\mathbf{g}_{lk}^T \mathbf{w}_{li}|^2 + \frac{\sigma_n^2 MK}{P_{\mathrm{BS}}}}. \tag{2.12}$$

Reverse-Link Data Transmission

The reverse link of a MU-MIMO setup is a multiple access channel, since all terminals transmit data over co-channelly. The BS separates the combined signal from all terminals by their spatial signatures. In case of synchronized multi-cell transmission, the signals from other cells contribute to the interference at every BS. The signal received at j^{th} BS is given by

$$\mathbf{x}_j = \sqrt{P_{\mathrm{T}}} \mathbf{G}_j \mathbf{a}_j + \sqrt{P_{\mathrm{T}}} \sum_{l=1, l \neq j}^{L} \mathbf{G}_{jl} \mathbf{a}_l + \mathbf{n}_j, \tag{2.13}$$

where P_{T} is the transmit power assumed equal for all the terminals, $\mathbf{a}_l$ is the vector of symbols transmitted by terminals in the l_{th} cell, and n_j is the thermal noise at the BS. The first term in Eq. 2.13 models the composite signal from all terminals received at the BS, and the second term corresponds to the inter-cell interference.

The BS processes the received signal to maximize the SNR over each link (maximum ratio combining, MRC) or reject the interference from other links in the cell (zero forcing, ZF receiver). The formulations for these receivers is given by Eq. 2.14 and Eq. 2.15 respectively.

$$\mathbf{Z}_j^{MRC} = \frac{\widehat{\mathbf{G}}_j^{H}}{\|\widehat{\mathbf{G}}_j\|}, \tag{2.14}$$

$$\mathbf{Z}_j^{ZF} = \frac{\widehat{\mathbf{G}}_j^{\dagger T}}{\|\widehat{\mathbf{G}}_j^{\dagger}\|}. \tag{2.15}$$

The advantages and drawbacks of each technique are analogous to the forward-link transmission. The schematic for reverse-link data transmission is shown in Fig. 2.5. It is worth noting here that the receiver noise at the BS is usually less compared to the terminals, because of higher precision components and better control over the ambient temperature. In this thesis, we focus on the forward-link data transmission in terms of effective SINR and throughput. The extension to reverse-link is straightforward and left for further work.

Figure 2.5: Reverse-link data transmission and decoding in MU-MIMO

Note that we have assumed perfect reciprocity of the channel in the preceding discussion of forward- and reverse- link data transmission. The propagation channel fulfills reciprocity in the absence of scatterers with unusual magnetic properties. However, the receiver chains at the BS and the terminals are not reciprocal in general, and have to be calibrated to satisfy this condition.

In addition to the linear processing techniques discussed above, there exist non-linear methods like successive interference cancellation, dirty paper coding, etc. that have been shown to improve the performance for small MU-MIMO systems at the cost of additional complexity.

2.3   Massive MIMO

The MU-MIMO technology has been incorporated in LTE Release 8 and further standards, with a maximum of 8 BS antennas expected to serve roughly an equal number of terminals with FDD operation [13]. The overhead of channel estimation in such a system grows rapidly with the number of BS antennas, making it non-scalable. A recent proposal in MU-MIMO research is the deployment of a large excess of BS antennas as compared to the number of terminals served, termed "massive MIMO" [14]. This technique promises spectacular gains in spectral and radiated energy efficiency over the traditional point-to-point MIMO systems. Under TDD operation and the assumption of channel reciprocity, the BS can estimate forward- and reverse-link channels using reverse-link pilot sequences. The number of orthogonal pilot resources scales as the number of terminals served, and is independent of the BS array size. Extra BS antennas are cost-free in terms of pilot resources, and have been shown to always improve the channel estimates in a single-cell setup with linear MMSE estimation [2]. However in a multi-cell setup, the pilot sequences need to be reused by terminals in nearby cells and potentially contaminate the channel estimates.

The BS in a massive MIMO system exploits the large number of degrees of freedom to focus sharp beams towards the terminals using simple linear processing. As the size of BS array increases without bound, the channel vectors for the terminals become asymptotically orthogonal. Optimal performance can then be

achieved with maximum ratio transmission/combining in the forward/reverse links respectively. The narrow beams to each terminal suppresses intra-cell interference, so the required transmit energy per bit vanishes. The dominant impairment to communication is the inter-cell contamination of channel estimates during pilot-based training [14].

A practical massive MIMO system is expected to have up to several hundreds of BS antennas serving a few tens of terminals [15]. The envisioned of the BS is to spatially multiplex the terminals over all available time-frequency resources, increasing the spectral efficiency by an order of magnitude over current systems. The burden of processing complexity lies with the BS that allows high data rates to be achieved even for cheap, single-antenna terminals. The universal time and frequency reuse reduces the overhead of scheduling, simplifying the multiple-access layer. At the same time, the energy efficiency increases by an order or more of magnitude. The low transmit power at the BS antennas can be served by hundreds of low-cost linear amplifiers instead of the traditional high-cost, ultra-linear amplifiers operating at relatively high output powers. The additional degrees of freedom at BS can be used to place nulls in the direction of interferers, improving the SINR.

The massive MIMO system faces several challenges in its implementation, though the solutions for most of them are available at least in theory. The majority of current LTE systems rely on FDD operation, which is intractable from the perspective of massive MIMO. Even with TDD operation, channel reciprocity is a simplifying assumption that is not expected to hold in practice. However, calibration methods exist that can make the forward and reverse links approximately reciprocal and facilitate massive MIMO. The gains with MRT/MRC depend on the orthogonality of channel responses to different terminals, which set the conditions for *favorable propagation*. Recent studies point out that favorable propagation might be achievable even for moderate antenna sizes in practice [16]. The resources available for channel estimation are fundamentally limited by the coherence interval. In a multi-cell setup, these resources have to be reused across cells, leading to pilot contamination. The impact of pilot contamination on real systems is hotly debated, and is the topic of this thesis.

Pilot-Based Channel Estimation

In a MU-MIMO setup, the BS jointly precodes data for all served terminals in the forward link, and jointly decodes the reverse link. The BS relies on good channel knowledge to create spatially multiplexed data streams for each terminal for joint encoding/decoding. In Sec. 3.1, we introduce the concept of pilot symbols and note their use in context of current LTE systems. We study frequency- and time-division duplex (FDD and TDD) operation in the context of the LTE frame structure, and its impact on the overhead of channel estimation. Next, in Sec. 3.2 we study common channel estimation techniques while employing reverse-link pilot sequences. We investigate the LS and MMSE approaches in a multicell MU-MIMO setup and develop metrics for comparing their performance. In Sec. 3.3, we introduce the problem of *pilot contamination* and its impact on forward-link MU-MIMO transmission. We derive the expression for achievable rate for forward-link transmission assuming a finite channel coherence interval, while taking into consideration the overhead of channel estimation.

3.1   Channel Training Using Pilot Symbols

A common technique for channel estimation is to transmit a known sequence of symbols (pilots) and evaluate the effect of the channel on these symbols at the receiver [17]. In LTE, these symbols are known as *reference symbols*. In the forward link, the BS periodically transmits cell-specific reference signals (CRS) that are used by terminals for initial acquisition, CQI measurement and channel estimation for coherent detection. These CQI values may be reported to the BS over a control channel which uses it for closed-loop power control, rate adaptation and scheduling [6]. The terminals may transmit reference signals as well, for coherent demodulation of the reverse link at the BS. In case of MU-MIMO, the BS requires fairly accurate channel estimates for coherent processing of the forward and reverse links.

FDD vs TDD operation

The forward- and reverse-link transmissions in current LTE systems are duplexed over time (TDD) or frequency (FDD) to minimize self-interference[1]. In TDD mode, the links occupy the same frequency band but are assigned different slots in time, while in FDD they occupy non-overlapping frequency bands that allows simultaneous transmission and reception at the BS.

In the reverse link of both TDD and FDD, each terminals transmit pilot symbols that is orthogonal to other terminals within the cell to minimize intra-cell interference. The channel between a terminal and each BS antenna is estimated using the received pilot symbols from the corresponding terminal. Therefore, the required pilot resources depends on the number of terminals, but is independent of the number of BS antennas. This is particularly attractive for massive MIMO systems, where the BS array size is expected to be an order of magnitude larger than the number of terminals ($M >> K$). Moreover, TDD offers a distinct advantage: the estimates in the reverse link are valid for the forward link as well, albeit with some calibration [19]. However in FDD, the BS has to obtain forward link estimates separately, that increases the channel estimation overhead.
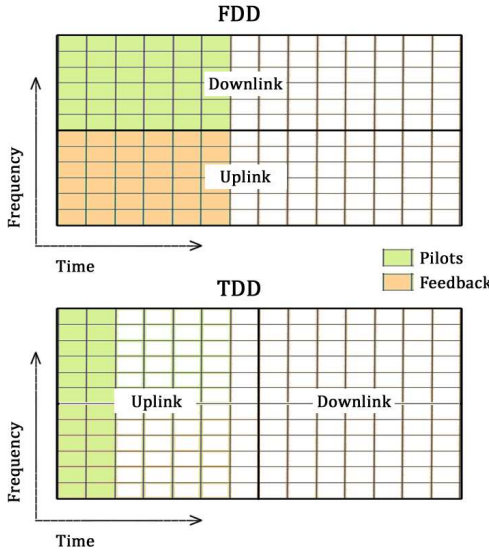
Figure 3.1: Forward-link channel estimation overhead for FDD and TDD operation modes

The forward-link channel may be estimated at the terminals using mutually orthogonal pilot transmissions from each BS antenna. The number of required

---

[1]Recently, a so-called "full duplex" technique for simultaneous transmission and reception over the same frequency band has been proposed and demonstrated, though it is still far from standardization [18]

time-frequency indices is therefore proportional to M. Each terminal reports the $M \times 1$ channel vector over the (spatially multiplexed) reverse link, consuming additional resources proportional to M. The estimation and feedback adds latency before the BS can use them for multiplexing the forward-link data, which scales proportional to M. Therefore in a massive MIMO system, FDD is generally infeasible using the pilot-based training schemes. Some alternative methods have been proposed that estimate the forward link from the eigenvalues of the reverse-link, which are untested in practical systems as yet [20][21].

3.2 Channel Estimation in the Reverse Link

We assume that the k^{th} terminal in every j^{th} cell, $j \in \mathcal{J}$ synchronously transmits a pilot sequence $\mathbf{s}_k$ comprising τ pilot symbols,

$$
\begin{aligned}
\mathbf{s}_k &= [s_{k1} \ s_{k2} \ ... \ s_{k\tau}]^T, && (3.1) \\
\sum_i \|\mathbf{s}_{ki}\| &= \tau P_{\mathrm{T}}.
\end{aligned}
$$

The $M \times \tau$ signal received at the j^{th} BS due to transmission of the k^{th} pilot sequence is given by

$$
\mathbf{X}_{jk} = \sum_{l \in \mathcal{J}} \mathbf{g}_{jkl} \mathbf{s}_k^T + \mathbf{N}_j, \tag{3.2}
$$

where $\mathbf{g}_{jkl}$ is the $M \times 1$ channel vector between j^{th} BS and the k^{th} terminal in l^{th} cell, and $\mathbf{N}_j$ is thermal noise at the BS, modeled as white Gaussian. Since we assume the pilot signals within a cell to be mutually orthogonal, we can safely ignore the reception of all other pilot sequences $\mathbf{s}_i, i \neq k$ at the BS during this analysis.

3.2.1 LS Estimation

The LS approach to channel estimation seeks to minimize the squared error between the received pilot sequence and its noise-and-interference free version [22]. The LS estimator determines a channel estimate, $\widehat{\mathbf{g}}_{jk}^{\mathrm{LS}}$ such that its distance from the actual channel is minimized,

$$
\widehat{\mathbf{g}}_{jk}^{\mathrm{LS}} = \arg\min_{\mathbf{g}} \|\mathbf{X}_{jk} - \mathbf{g}\mathbf{s}_k^T\|^2. \tag{3.3}
$$

$$
\begin{aligned}
\text{Let } \mathcal{E} &= \|\mathbf{X}_{jk} - \mathbf{g}\mathbf{s}_{jk}^T\|^2 \\
&\triangleq (\mathbf{X}_{jk} - \mathbf{g}\mathbf{s}_k^T)^H (\mathbf{X}_{jk} - \mathbf{g}\mathbf{s}_k^T) \\
&= (\mathbf{X}_{jk}^H \mathbf{X}_{jk} - \mathbf{s}_k^* \mathbf{g}^H \mathbf{X}_k - \mathbf{X}_{jk}^H \mathbf{g}\mathbf{s}_k^T + \mathbf{s}_k^* \mathbf{g}^H \mathbf{g}\mathbf{s}_k^T), && (3.4)
\end{aligned}
$$

where we have utilized the following vector properties:

$$
\begin{aligned}
(\mathbf{a} + \mathbf{b})^H &= \mathbf{a}^H + \mathbf{b}^H, \\
\mathbf{a}\mathbf{b}^H &= \mathbf{b}^H \mathbf{a}^H, \\
(\mathbf{a}^T)^H &= \mathbf{a}^*.
\end{aligned}
$$

To obtain the LS channel estimate, we can minimize $\mathcal{E}$ by setting its partial derivative with respect to $\mathbf{g}^H$ to 0. The received signal and the channel are assumed to be deterministic and hence constant with respect to $\mathbf{g}^H$.

$$\frac{\partial \mathcal{E}}{\partial \mathbf{g}^H} = 0$$

Utilizing the value of $\mathcal{E}$ from Eq. 3.4 and taking partial derivative, we get

$$\mathbf{X}_{jk}\mathbf{s}_k^* - \mathbf{g}\mathbf{s}_k^T\mathbf{s}_k^* = 0. \tag{3.5}$$

$$\widehat{\mathbf{g}}_{jk}^{\text{LS}} = \mathbf{X}_{jk}\mathbf{s}_k^*(\mathbf{s}_k^T\mathbf{s}_k^*)^{-1}$$

$$= \mathbf{g}_{jk} + \sum_{l \in \mathcal{J}, l \neq j} \mathbf{g}_{jkl} + \frac{\mathbf{N}_j\mathbf{s}_k^*}{\tau P_{\text{T}}}, \tag{3.6}$$

where $\widehat{\mathbf{g}}_{jk}^{\text{LS}}$ is the LS channel estimate for the k^{th} terminal in the j^{th} cell.

The LS estimator has low complexity and treats the channel coefficients as a deterministic variable to obtain a "best-fit" estimate from the observed pilot signal. It makes no prior assumption about the channel statistics, and due to its simple implementation it the most common approach to channel estimation in practice. Another class of estimators exist that treat the channel as a stochastic variable, and exploit some prior knowledge of the channel statistics. These so-called *Bayesian* estimators have a higher complexity, but potentially perform better than the LS estimator in terms of the MSE of channel estimation. We study one particularly common Bayesian estimator, the MMSE estimator, in the next section.

MMSE Estimation

In Sec. 2.1, we described the spatial correlation of a channel across an antenna array. The spatial information of a terminal, contained in its channel covariance matrix, varies slowly compared to the small-scale fading. The MMSE estimator exploits the prior knowledge of covariance matrices to improve the channel estimates, by amplifying the signal from spatial direction of desired terminal and attenuating the interferers [22]. We derive the MMSE estimator in Appendix 7.1, and just present the final expression here.

$$\widehat{\mathbf{g}}_{jk}^{MMSE} = \mathbf{R}_{jk}\left(\sigma_n^2\mathbf{I}_M + \tau \sum_{l \in \mathcal{J}} \mathbf{R}_{jkl}\right)^{-1} \mathbf{S}_k^H\mathbf{x}_{jk}. \tag{3.7}$$

$$\text{where } \mathbf{x}_{jk} = \text{vec}(\mathbf{X}_{jk}),$$

$$\mathbf{S}_k = \mathbf{s}_k \otimes \mathbf{I}_M,$$

and $\mathbf{R}_{jk}$ is the $M \times M$ covariance matrix of the channel between the j^{th} BS and the k^{th} terminal in corresponding cell. $\mathbf{R}_{jkl}$ is the covariance matrix of the channel between the j^{th} BS and the k^{th} terminal in the l^{th} cell, and $\otimes$ denotes Kronecker product. The rank of the covariance matrices depends upon the angular spread of corresponding channel. A full-rank covariance matrix corresponds to a wide

angular spread, implying that the array cannot effectively resolve the angular location of the transmitter. On the other hand, a low rank indicates narrow angular spread. The accuracy of the MMSE estimator depends on the overlap of the angular spread of interfering terminals. In case of large angular spread and several terminals, there may be significant overlap and the performance of MMSE estimator suffers. However if the angular spreads do not overlap significantly, the MMSE estimator can be shown to eliminate the pilot interference almost completely [23].

We observe that theoretically, the MMSE estimator has significantly higher implementation and processing complexity than the LS estimator. It requires the knowledge of all cross-channel covariance matrices at all BSs, that must be estimated prior to MMSE channel estimation. In practice, this imposes significant overhead and additional latency on the system. To reduce this overhead, we can assume that the interference from terminals located more than a few cells away to be negligible and not estimate the corresponding cross-channel matrices, at the cost of slightly poorer estimator performance. In terms of processing requirements, the complexity of $M \times M$ matrix inversion required during evaluation of MMSE estimate is proportional to the cube of array size, and may be especially problematic for massive MIMO systems.

Normalized MSE

In Ch. 5, we evaluate the performance of the LS and MMSE estimators in terms of the squared error of estimated channel coefficients. A useful metric for this analysis is the MSE of channel estimates, where the expectation taken over several channel realizations,

$$\mathcal{M}_{jk} \triangleq E\{\|\widehat{\mathbf{g}}_{jk} - \mathbf{g}_{jk}\|^2\}. \tag{3.8}$$

On the one hand as the array size increases, the MSE increases because of additional channel coefficients that have to be estimated. On the other hand, additional antennas improve the processing at the BS for forward and reverse links, even with imperfect channel estimates. It is therefore useful to normalize the MSE with the BS array size. Moreover, we are interested in the estimation error as a fraction of the channel gain - a small error for a poor channel is more significant than a relatively larger error for a strong channel. Combining these two requirements, we normalize the MSE with respect to expected channel gain and evaluate the estimators in terms of the normalized MSE (NMSE), η_{jk} defined as

$$\eta_{jk} \triangleq \frac{E\{\|\widehat{\mathbf{g}}_{jk} - \mathbf{g}_{jk}\|^2\}}{E\{\|\mathbf{g}_{jk}\|^2\}}. \tag{3.9}$$

3.3   Pilot Contamination

As discussed in Sec. 2.1, the channel estimates obtained above are useful only within the coherence interval, after which the channel must be estimated again. Moreover, the maximum number of mutually orthogonal pilot sequences is fundamentally limited by τ which must be smaller than number of coherent time-frequency elements, $\tau \leq N_{\text{coh}}$. Within a cell, the K terminals always use orthog-

onal pilot sequences to eliminate intra-cell pilot interference ($\tau \geq K$). However depending on the value of N_{coh}, these sequences may have to be reused in other cells, which leads to inter-cell interference. In Eq. 3.2, we observed the effect of this interference on the received pilot signal, known as the "pilot contamination" effect (Fig. 3.2). In particular, we found that the received pilot signal is contaminated with the transmissions from terminals in $l \in \mathcal{J}, l \neq j$ cells reusing the same sequence. The worst-case pilot contamination occurs when each cell reuses the same set of mutually orthogonal pilot sequences.

Figure 3.2: Contamination of reverse-link pilot signal due to reuse of pilot sequences in other cells

During forward-link transmission, pilot contamination adversely affects the precoding weight vectors for the k^{th} terminal. In case of MRT precoding discussed in Sec. 2.2, the BS uses the channel estimates to focus signal energy at the location of the k^{th} terminal. Due to the contaminated channel estimates, some of the transmit power *leaks* towards the k^{th} terminals in other cells belonging to $\mathcal{J}$, attenuating the signal power at the desired terminal. Moreover, this leakage appears as focused interference at the these terminals, decreasing their SINR (Fig. 3.3). Thus, the impact of pilot contamination on forward-link data transmission with MRT precoding is twofold: the signal power lost for desired terminal appears as sharply focused interference at other terminals that reuse the same pilot sequence!
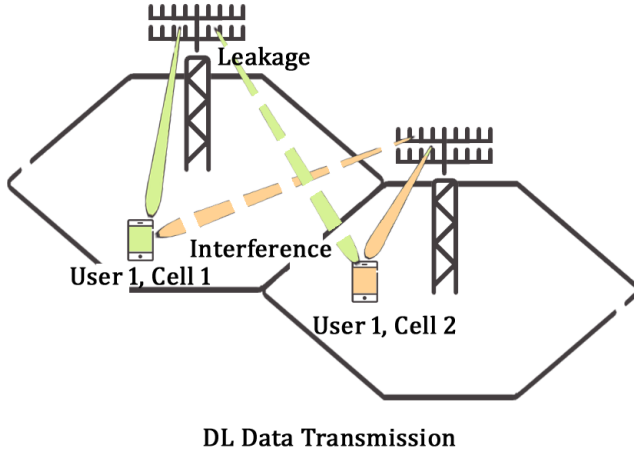
Figure 3.3: Effect of pilot contamination on forward-link data transmission with MRT precoding

Achievable Forward-link Rate

We have derived the forward-link SINR $\Gamma_{jk}^{\mathrm{FL}}$ for the k^{th} terminal of the j^{th} cell in (2.12), assuming TDD operation and perfect channel reciprocity. Here, we derive the achievable per unit bandwidth over the LTE forward-link grid. We assume Gaussian data symbols and interference, which implies a rate proportional to the logarithm of $1 + \Gamma_{jk}^{\mathrm{FL}}$. The OFDM symbol duration T_{s} in LTE consists of the data symbol and a cyclic prefix T_{CP} that protects against inter-symbol interference. Assuming N_{coh} coherent time-frequency resources, we allow N_{pilot} resources for channel estimation. The rate can then be obtained as

$$\mathcal{R}_{jk} \quad = \quad \left(\frac{N_{\mathrm{coh}} - N_{\mathrm{pilot}}}{N_{\mathrm{coh}}} \right) \left(\frac{T_{\mathrm{s}} - T_{\mathrm{CP}}}{T_{\mathrm{s}}} \right) \log_2(1 + \Gamma_{jk}^{\mathrm{FL}}) \; bits/s/Hz, \quad (3.10)$$

$$\mathcal{R}_{j} \quad = \quad \sum_{k=1}^{K} \mathcal{R}_{jk}, \quad (3.11)$$

$$\bar{\mathcal{R}} \quad = \quad \frac{\sum_{l=1}^{L} \mathcal{R}_{j}}{L}, \quad (3.12)$$

where $\mathcal{R}_{jk}$ is the forward-link achievable rate for k^{th} terminal in j^{th} cell, $\bar{\mathcal{R}}_{j}$ is the sum rate for j^{th} cell and $\bar{\mathcal{R}}$ is the per-cell sum rate of the network.

Techniques to Mitigate Pilot Contamination

In a massive MIMO system, the pilot resources are designed to be orthogonal within a cell, but are reused by terminals in other cells. The reuse of pilots in nearby cells causes interference during the channel estimation phase, known as *pilot contamination* and has been studied in [24]. We discussed this problem in detail in Ch. 3, and it has been investigated in the context of massive MIMO systems in [25]. In this chapter, we study techniques to mitigate pilot contamination in massive MIMO systems and analyze the scenarios under which they are effective in improving the channel estimation.

4.1 Pilot Open-Loop Power Control

The average received power of a pilot signal at the BS depends on the pathloss and transmit power of the mobile terminal. Furthermore, a terminal located closer to the BS usually has a lower pathloss, and hence a higher signal power at the BS. The terminals at cell-edge usually have a higher pathloss and therefore a low signal power even while transmitting at their maximum power. This makes them susceptible to interference from non-orthogonal pilot transmissions in neighboring cells. On the other hand, terminals close to the BS enjoy a reduced pathloss, resulting in a better SINR. The impact of interference on reverse-link pilot signals is shown in Fig. 4.1.

   Transmit power control in the reverse link, commonly known as uplink power control, is an important technique for the management of shared radio resources. We study a pilot open loop power control (pilot OLPC) scheme that allows the terminal to adjust the transmit power of its pilot signal based on its estimate of the pathloss to its serving BS. The terminal estimates pathloss from forward-link reference signal and tries to match a network-wide desired pilot signal strength at the BS [26]. This technique balances the received pilot SNR at the serving BS of a terminal with the interference generated at neighboring BSs. Pilot OLPC attempts to maximize SNR fairness of pilot signals for cell-edge terminals by compensating for their additional pathloss. This *pathloss compensation* has the intended effect of reducing overall pilot interference in the network at the cost of SNR degradation for terminals located close to the BS.

   Power control in LTE is a combination of open- and closed-loop techniques [27]. OLPC provides a coarse operating point for the terminal power in terms of

transmit power per RB. With closed-loop power control, the BS sends periodic power increment/decrement messages based on the SINR estimates, that defines a tighter short-term operating point. Power control is dependent only on the large-scale fading for the corresponding terminal, and hence is quite simple to implement. The operating points and compensation factor rely on existing and standardized measurements. Here, we ignore closed-loop power control and focus on the transmit power set by OLPC. The open-loop operating point for transmit power per RB depends on two factors: (i) a semi-static base level, P_0, assumed to be same across the network, and (ii) an open-loop pathloss compensation component. The k^{th} terminal in the j^{th} cell estimates forward-link pathloss $\mathcal{L}_{jk}$ and sets its transmit power P_{jk} according to the formula

$$P_{jk} \quad = \quad \min(P_0 + \alpha \cdot \mathcal{L}_{jk}, P_{\max}) + 10 \log_{10} M,$$

where α is the pathloss compensation factor, P_{max} is the maximum terminal power per RB and M is the number of allocated RBs. $\alpha = 1$ implies full pathloss compensation, ensuring maximum fairness for the cell edge terminals. The value of P_0 depends on the SNR target at the BS for channel estimation, usually set between 5 and 15 dB.

In presence of OLPC, the terminals with strong channel conditions at the BS back-off their transmit power reduce their SNR to the desired level. In the process, they reduce their interference power at all nearby BSs by the same factor. The cell-edge terminals sharing the same pilot sequence continue to transmit at a high power, improving their SINR as described in Fig. 4.1a. On the contrary, SINR for a terminal located close to the BS degrades because of the attenuation of its signal power with OLPC. The SINR of pilot signal depends with the accuracy of channel estimates at the BS. With 'good enough' channel estimates, the BS can improve the per-cell sum rate using spatial beamforming in the reverse- and forward-links. The gains in rate in a multi-cell setup with OLPC depend on the operating point P_0, the number of terminals in each cell and their distribution within the cells.

For our simulations, we assume that the BS and the terminals have perfect knowledge of the pathloss and can accurately, independently determine the pilot transmit power with OLPC. In practice, the terminal averages the pathloss estimates from forward-link pilots over 100-200 ms to average out the small scale fading and calculate its pilot transmit power, which is periodically signaled to the BS via uplink control channels for channel estimation. We assume the criteria for goodness of channel estimates for spatial multiplexing as NMSE $\leq$ 0dB. At NMSE = 0dB, the mean square error of channel estimation is equal to the average channel gain, which implies that half of the transmitted signal power with maximum ratio transmission is expected to be directed away from the desired terminal.

(a) Pilot reuse at cell-edge and close to BS

(b) Pilot reuse by close cell-edge terminals

(c) Pilot reuse close to serving BSs

Figure 4.1: Reverse-link pilot SINR with all terminals transmitting at maximum permissible power

4.2 Less Aggressive Pilot Reuse

Full pilot reuse leads to maximum inter-cell interference during channel estimation, which can be mitigated using a less aggressive pilot reuse factor. Pilot reuse is analogous to the traditional frequency reuse in the sense that terminals within the pilot reuse area can utilize only a fraction of the time-frequency resources, during the channel estimation phase. However with pilot reuse, each terminal is free to use all the available resources for communication for the rest of the coherence interval. The *pilot reuse factor* $1/U$ is the rate at which pilot resources may be reused in the network, where U is the number of cells that are assigned orthogonal pilots. A factor $U > 1$ always reduces the pilot contamination effect by assigning orthogonal pilots to neighboring cells, the next-neighbor cells and so on. The total number of unique time-frequency elements reserved for pilot transmission are KU, where K is the number of terminals per cell.

The trivial case of pilot reuse is full pilot reuse with $U = 1$. In this case, we reserve K time-frequency indices for pilot sequences as shown in Fig. 4.2a for $K = 12$. The indices might be located anywhere within the resource block without loss of generality. We generate K orthogonal pilot sequences spanning these indices, that are distributed at random or algorithmically among the terminals within each cell. Since the cells are assumed to be synchronized, the pilot transmissions are synchronized across cells as well.

(a) Pilot reuse 1

(b) Pilot reuse 1/3

(c) Pilot reuse 1/3 pattern
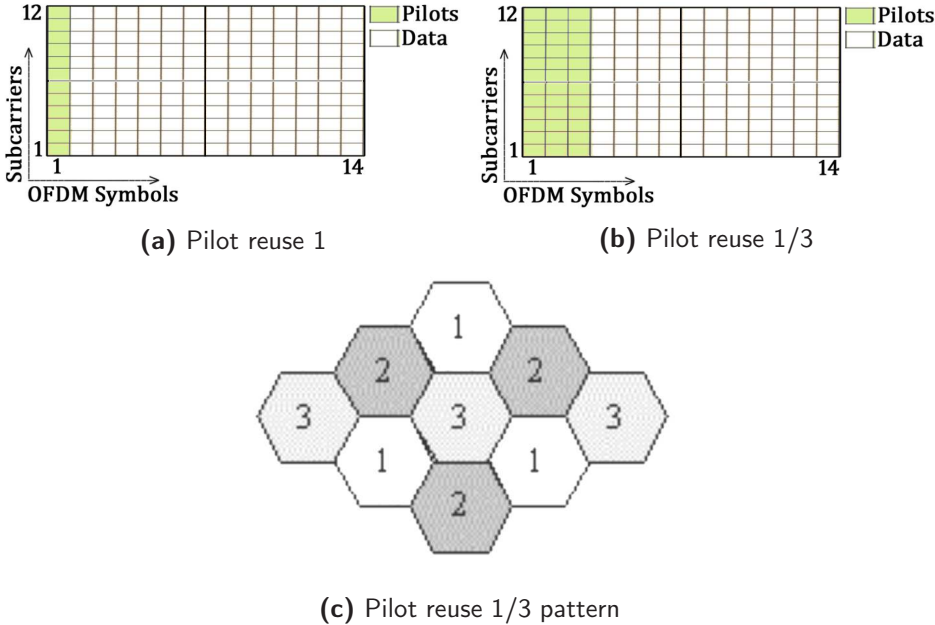
Figure 4.2: Time-frequency resources spent on pilots for $K = 12$ with reuse 1 and 1/3 respectively

Here we consider a less aggressive reuse factor in detail, where orthogonal pilot sequences are assigned to interfering terminals in the adjacent cells. In case

of hexagonal network layout, the smallest reuse factor that ensures orthogonal pilots in adjacent cells is 1/3 using a reuse pattern as described in Fig. 4.2c. To implement this reuse factor, we require $3K$ time-frequency resources to generate 3 groups of K mutually orthogonal pilot sequences each. We allocate a *pilot group* to each cell according to the reuse pattern, and the pilots within that group are distributed to the terminals as in the case of full pilot reuse. Similarly, we can define pilot reuse schemes for $U = 4, 7, 9$ etc. that successively reduce the pilot contamination. A larger U implies a larger distance r between cells using the same set of pilot sequences.

Clearly, a higher value of U improves the channel estimates but at the cost of an increased overhead of pilot resources, limiting the resources available for data within a coherence interval. The improved estimates are used for better forward- and reverse-link beamforming at the BS to serve multiple terminals via spatial multiplexing. Thus, there exists a trade-off between the accuracy of MU beamforming and the resources available for spatially multiplexed data. It is worth noting here that since received interference power varies inversely with $r^{\gamma}, \gamma > 1$, there is smaller successive gain in terms of pilot contamination as we increase U. Instead, a larger pool of pilot sequences might be utilized for serving additional terminals within each cell with a conservative pilot reuse factor, at the cost of poorer spatial multiplexing.

Soft Pilot Reuse

The received power in the reverse link depends on the transmit power of the terminal and also its location with respect to the BS. A cell-edge terminal usually has poor SNR at its serving BS and causes higher interference at nearby cells, than a terminal located far from the cell edge. If two cell-edge terminals from different cells are located close to each other and share the same pilot sequence, they suffer from significant pilot contamination (Fig. 4.1b). On the contrary, terminals located close to the BS are robust against interference from other cells (Fig. 4.1a, 4.1c). The pilot reuse scheme described above, though attractive on account of simple implementation, gives equal priority to all terminals while assigning pilot sequences. We develop a *soft pilot reuse* (SPR) scheme that mitigates pilot contamination by assigning additional orthogonal pilot resources to cell-edge terminals.

We reserve a fraction of pilot resources for the cell-edge terminals in all cells, that are assigned with a less aggressive reuse factor $1/U_{\mathrm{E}}$. The terminals close to serving BS are expected to cause less interference and follow a reuse factor $1/U, U < U_{\mathrm{E}}$. The BS classifies ρK terminals with the lowest SNR as being at the cell edge, where $\rho \leq 1$ is a network-wide parameter. The remaining terminals are assumed to be closer to the BS and less sensitive to inter-cell interference. The total pilot resources spent on cell-edge terminals is therefore $\rho K \cdot U_{\mathrm{E}}$. Similarly, the number of resources utilized for pilot transmission by the rest of the terminals is $(1 - \rho)K \cdot U$.

We consider the situation when half the terminals are classified as cell-edge terminals ($\rho = 0.5$) that reuse pilots with $U_{\mathrm{E}} = 3$. The other half of the terminals are assumed better insulated from inter-cell interference and have a pilot reuse

factor $U = 1$. The reuse pattern is shown in Fig. 4.3. The total number of time-frequency indices spent on pilots is therefore $0.5K \cdot 3 + (1 - 0.5)K \cdot 1 = 2K$, so this scheme can also be described as "1/2 SPR". In contrast, pilot reuse 1 and 3 involve an overhead of K and $3K$ coherent resources respectively. During simulations, we evaluate the per-cell sum rate to benchmark SPR against these (relatively) simpler schemes while accounting for the pilot overhead.

Figure 4.3: Distribution of pilot sequences with soft pilot reuse (SPR)

We assume that the BS has perfect knowledge of reverse-link SNR for each terminal it serves. In practice, the terminal estimates the forward-link pathloss and reports it along with its transmit power to the BS over a control channel. This adds some amount of inaccuracy and latency during pilot assignment. Moreover since the terminal positions change with time, the pilots need to be reassigned based on updated SNR information. These overhead make the SPR scheme less tractable than simple pilot reuse. However with short coherence intervals or for a large number of terminals, the pilot resources can become quite expensive, and SPR can be of advantage. With SPR, the pilot contamination between cell-edge terminals is effectively mitigated. As such, the terminals close to the BS might become dominant cause of interference, especially if they transmit at maximum power. In this case, OLPC can be incorporated to further reduce contamination as discussed in the previous section.

4.3 Inter-cell Coordinated Pilot Allocation

In an urban macro scenario, small scale fading causes significant channel variation with time in the order of a few milliseconds. However, a terminal's angular location with respect to the BS changes relatively slowly, in the order of a few seconds. The spatial information of the terminal is characterized by the second order statistics of its channel vector, known as the covariance matrices and introduced in Ch. 2. The MMSE estimator discussed in Ch. 3 exploits prior knowledge of channel covariance matrices to amplify the desired signal and attenuate the interference and noise. If

the terminals are well-separated in space such that the angular spread of incoming signals at the BS do not overlap, the MMSE estimator is effective at rejecting the interference signal with a large enough array. This motivates devising a technique to assign pilot sequences to spatially well-separated terminals.

A slow-rate, inter-cell coordinated pilot assignment (CPA) algorithm that investigates this approach has been proposed in [23]. CPA relies on all cross-channel covariance matrices being available at all BS. This information is used to evaluate a closed-form expression for the expected channel estimation error when a given set of terminals reuse a pilot sequence. A greedy algorithm iteratively assigns pilots in every cell to the terminal that minimizes this error. We now describe the CPA algorithm in detail.

We assume that K orthogonal pilot sequences have to be assigned to the terminals within a $L-$cell network with pilot reuse 1. The set of all terminals reusing the k^{th} pilot sequence is denoted by $\mathcal{U}_k$. We use the network utility function $F(\mathcal{U}_k)$ described in [23] as a measure of the pilot contamination,

$$F(\mathcal{U}_k) \quad = \quad \sum_{j=1}^{|\mathcal{U}_k|} \frac{\mathcal{M}_{jk}}{tr\{\mathbf{R}_{jk}\}}. \tag{4.1}$$

$$\mathcal{M}_{jk} \quad = \quad tr\left\{\mathbf{R}_{jk} - \mathbf{R}_{jk}^2\left(\frac{\sigma_n^2}{\tau}\mathbf{I}_M + \sum_{l=1}^{L}\mathbf{R}_{lk}\right)\right\}, \tag{4.2}$$

where $|\mathcal{U}_k|$ is the cardinal number of the set $\mathcal{U}_k$. $\mathcal{M}_{jk}$ is the expected MSE of channel estimation for the terminal assigned k^{th} pilot sequence in j^{th} cell. A higher value $F(\mathcal{U}_k)$ implies poor spatial separation of the terminals and consequently higher contamination using the MMSE estimator. We utilize this property to assign the pilots in each cell as described in Algorithm 1

Algorithm 1 CPA Algorithm

---

    for $k = 1, 2, ..., K$ **do**
      $\mathcal{U}_k \leftarrow \emptyset$
      for $j = 1, 2, ..., L$ **do**
        $\kappa_j = \arg\min_{k \in \mathcal{G}} F(\mathcal{U} \cup \{\kappa_j\})$
        $\mathcal{U}_j \leftarrow \mathcal{U}_j \cup \{\kappa_j\}$
      end for
    end for

Channel covariance matrices might be obtained in practice by taking expectation of the outer product of the channel estimate with its conjugate transpose over several time instants. To obtain good statistics, accurate channel information is required for tens or even hundreds of time instants. This leads to a significantly higher overhead for the MMSE estimator as compared to the simple LS estimator. Recently, methods that multiplex pilots with data have been shown to be efficient in estimating covariance matrices with relatively low overhead. However, CPA requires the sharing of all covariance matrices across the network with a slow rate, adding to the overhead in the backhaul. Under the constraint of this overhead, the

algorithm has been shown to completely eliminate interference, but only for the initial terminal in a two-cell network in [23]. We evaluate CPA against random allocation of pilots and MMSE estimation in a 21-cell network in Ch. 5.

Simulation Setup and Numerical Results

In this chapter, we evaluate the pilot contamination effect and the performance of the proposed mitigation schemes in a massive MIMO system. In Sec. 5.1, we describe the simulation setup and explain the choice of parameters. This setup is used for the numerical evaluations in Sec. 5.2. First, we investigate the impact of pilot contamination on channel estimation NMSE in case of full pilot reuse ($U = 1$), when LS and MMSE estimators are used. Next, we evaluate the effect of pilot OLPC on the NMSE at each terminal.

We then use the pilot-contaminated channel estimates for MRT and ZF precoding. We evaluate the forward-link terminal SINR and the per-cell sum rate ($\bar{\mathcal{R}}$), and compare it against perfect channel knowledge at the BS. Later in this section, we evaluate the mitigation techniques discussed in Ch. 4 within our setup, and discuss the results. A summary of evaluated techniques to analyze and mitigate the impact of pilot contamination, and the performance metrics used, is presented in Table 5.1. We evaluate the SINR and $\bar{\mathcal{R}}$ with MRT as well as ZF precoding for all simulated scenarios.

	$U = 1$	$U = 3$	$U = 1/2$ SPR	$U = 1$, CPA
NMSE	Fig. 5.2, 5.3b	5.9a	-	5.14
SINR	Fig. 5.5a	5.10a, 5.10b	5.13a	-
$\bar{\mathcal{R}}$	Fig. 5.6, 5.8	5.11a, 5.11b	5.13b	5.15a, 5.15b

Table 5.1: Simulated techniques to analyze the effect of pilot contamination and mitigation techniques on channel estimation and forward-link transmission

5.1 Simulation Setup

We simulate the ITU Urban Macro model described in [7]. The simulation area comprises 7 three-cell sites, with hexagonally shaped cells. Each of the $L = 21$ cells is served by a time-synchronized BS with directional antenna array. The cells at the edge of simulation area experience smaller inter-cell interference, since they have less number of adjoining cells. Therefore, to approximate similar interference conditions at all cells, we incorporate a virtual wrap-around of cells while generating the channel matrices [28]. The terminals are distributed randomly within each cell except for an *exclusion zone* around the BS. The exclusion zone is used

to ensure that forward- and reverse-link communication happens in the far field
of transmitting antenna, and plane waves can be assumed at the receiver, for
which the channel model is valid. We assume that the terminals move in random
directions with vehicular speeds.
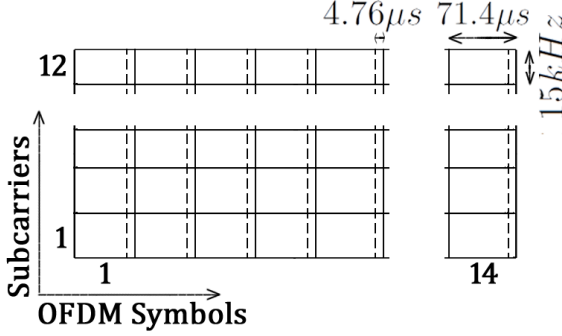
Figure 5.1: LTE forward-link RB pair

We assume that the channel matrix between a BS and the terminals is recipro-
cal in the forward and reverse links, which are duplexed in time (TDD operation).
All communication is using OFDM symbols over the LTE forward-link grid shown
in Fig. 5.1. This grid comprises OFDM symbols with useful symbol duration 66.7
μs and a guard interval of 4.76 μs. Each OFDM symbol is spread over 12 subcar-
riers spaced 15 kHz apart. We assume that the channel remains constant during
the entire coherence interval (block fading channel).

We now evaluate the coherence interval for our scenario. In Eq. 2.2, the path
length of MPCs is of the order of inter-site distance (500 m). We use conservative
value of $\Delta = 1500$, that leads to coherence bandwidth $B_c = 180$ kHz (12 subcar-
riers). In Eq. 2.3, the relative speed of a wavefront with respect to its receiver is
of the order of terminal speeds (60 kmph) since the BS is fixed. We assume con-
servative value of $v = 250$ kmph at the given carrier frequency of 2 GHz, leading
to a coherence time $T_c = 1$ ms (14 OFDM symbols). These dimensions of the
coherence interval (14 OFDM symbols $\times$ 12 subcarriers) corresponds to a LTE
RB pair [6].

We consider a single *drop* for each simulation run, implying that except fast
fading, all other random properties do not change during the simulation (e.g.
location of terminals, slow fading, angle of arrival etc) [7]. We obtain the channel
matrices between each BS and terminal at 40 time instants at intervals of 10 ms.
The channel coefficients are obtained at the same frequency, and are assumed to
be constant over the 12 subcarriers of a RB. We use these channel realizations to
estimate the covariance matrices in case of MMSE estimation.

Within a cell, all terminals are assigned an orthogonal pilot sequence, that
might be reused by terminals in other cells. Each pilot sequence is a vector con-
taining τ values that are stacked onto subcarriers of the OFDM symbols discussed
above. We obtain the k^{th} pilot sequence from the corresponding column of a $\tau \times \tau$
identity matrix. In practice, pilot sequences with good orthogonal separation and
lower peak-to-average (PAR) ratio may obtained from Zadoff-Chu sequences [6].

The impact of non-orthogonal pilot sequences on contamination of channel estimates is left for future investigation.

The links between the terminals and an elevated BS span a narrow range of elevation angles, but wide range of azimuth angles. Therefore, we choose BS antennas with wide horizontal beamwidth and narrow vertical beamwidth to serve the terminals within its cell. Additionally, we use a uniform linear array (ULA) geometry to allow good resolution in the azimuth [29]. The carrier wavelength for our operating frequency of 2 GHz is $\lambda_c = 15$ cm. The corresponding size of a 100-antenna ULA with $0.7\lambda_c$ spacing is about 10 m, much larger than existing BS deployments. In practice, such an array may be difficult to deploy and maintain. Instead, we can expect that the physical size of the array might be substantially reduced with a rectangular/semi-cylindrical geometry or by operating at higher frequencies.

We assume *full buffer* operation, i.e. all the terminals have content to transmit/receive at all times. This is a simplifying assumption, to focus on the effect of pilot contamination on forward-link transmission to few tens of active terminals. In practice, the traffic may be bursty, that allows the scheduling of additional terminals. The parameters we use are summarized in Table 5.2. The BS antenna elements are modeled according to specifications in [30], which may be deployed in practice using Kathrein antennas.

Scenario	ITU Urban Macro [7]
Network Deployment	21-cell hexagonal grid
Inter-site distance	500 m
Exclusion radius	35 m
Terminals per cell (K)	{3,6,**12**,24}
Terminal speed	60 kmph
BS transmit power (P_{BS})	0.067W per subcarrier
Max. terminal transmit power (P_{T})	23 dBm over 20 MHz
Carrier Frequency (f_{c})	2 GHz
Subcarrier spacing	15 KHz
BS array	100-antenna uniform linear array (ULA)
Tilt	11°
BS antenna	Vertically Polarized
BS antenna spacing	0.7 λ_c
BS max. antenna gain	18 dBi
BS 3dB horizontal beamwidth	65°
BS 3dB vertical beamwidth	6.5°
BS antenna noise figure	5 dB
Terminal antenna	Omnidirectional, Vertically Polarized
Terminal antenna noise figure	9 dB

**Table 5.2:** Simulation Parameters

5.2   Results and Discussion

Channel Estimation with Full Pilot Reuse

We evaluate the impact of pilot contamination on channel estimation NMSE when the same set of pilot sequences is reused ($U = 1$) in each of the $L = 21$ cells. In this case, the pilot sequences are columns of a $K \times K$ identity matrix. Each terminal within a cell is assigned a unique sequence randomly, eliminating intra-cell interference. The terminals transmit at maximum permissible power level. The BS uses LS/MMSE channel estimators described in Ch. 3 to estimate the $M \times K$ channel matrix from the received pilot signal.

In Fig. 5.2, we evaluate the NMSE for $K = 12$ terminals per cell. We observe that in case of LS estimator (green curve), the NMSE for a large number of termi-nals is higher than 0 dB. This happens for terminals with poor channel conditions at the BS (typically cell-edge terminals), when the interfering pilot signal is strong. In practice, these estimates are too noisy to be used for precoding at the BS, and the terminals face an outage.

The MMSE estimator has access to spatial information and pathloss for each terminal from the covariance matrices. It exploits this knowledge to reject inter-cell pilot interference to improve the channel estimates for desired terminal (red curves in Fig. 5.2). We observe that if a larger array is used ($M = 100$, solid red), the MMSE estimator performs better compared to a smaller array ($M = 10$, dashed red) because of increased spatial information. With the MMSE estimator, the NMSE for most terminals is less than 0 dB, allowing effective precoding and processing of data signals.
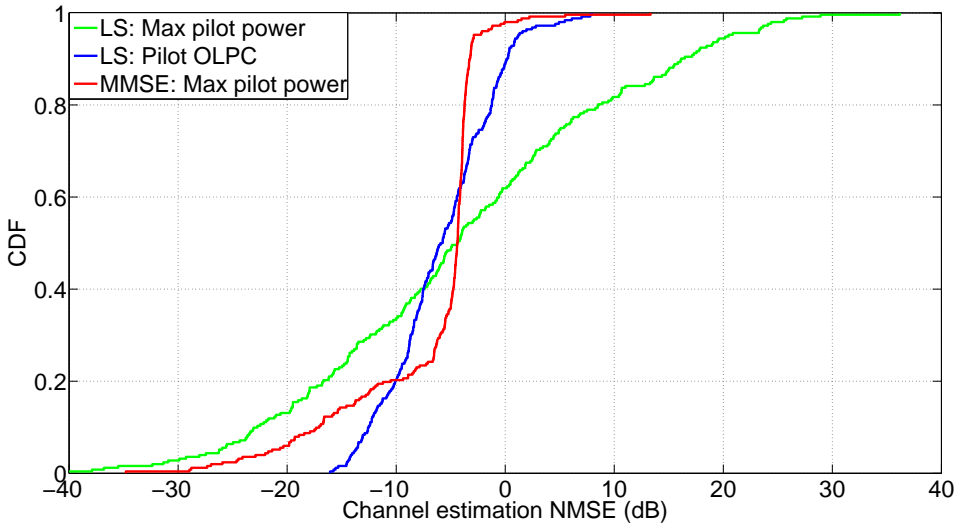

Figure 5.2: CDF of channel estimation NMSE for all terminals. LS estimator gives poor estimates for large number of terminals. Performance of MMSE estimator improves with M ($U = 1, K = 12$).

Next, we implement the pilot OLPC scheme described in Sec. 4.1 to back-off the transmit power of potentially strong interferers. We keep the desired SNR of received pilot signal at 10 dB and $\alpha = 1$ (full pathloss compensation). The revised pilot transmit powers with pilot OLPC are shown with a blue curve in Fig. 5.3a.
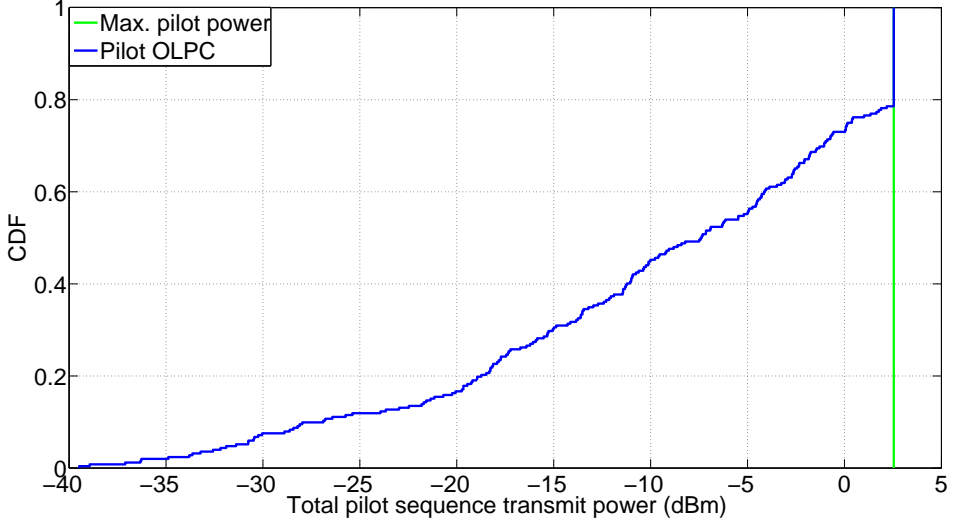
(a) Pilot OLPC: Terminals with strong channels back-off their transmit powers to improve SINR-fairness (desired pilot SNR at BS = 10 dB, $\alpha = 1$)

(b) CDF of channel estimation NMSE for all terminals. LS estimation with pilot OLPC (blue) balances the NMSE compared to maximum-power pilots (green).

Figure 5.3: Impact of pilot OLPC on LS channel estimation NMSE ($L = 21, M = 100, U = 1, K = 12$)

In Fig. 5.3b, we observe that the NMSE-fairness is improved with pilot OLPC (blue curve), and the most LS channel estimates are good enough for precoding (NMSE <0 dB). The corresponding CDF with MMSE channel estimation and maximum-power pilots is shown with a red curve for reference. In rest of the simulations, we always use pilot OLPC in case of full pilot reuse and LS estimator, to ensure that the system is within operating range.

A note about transmit power fairness: while comparing NMSE in presence and absence of pilot OLPC, we have to be aware of the difference in total power spent on pilot transmission. With pilot OLPC, the transmit power of several terminals is reduced, resulting in better energy efficiency and thus longer battery life. For fair comparison in case of max. power transmission, we should reduce all terminal powers such that the their sum becomes equal to that with pilot OLPC. However, since the received pilot signal is limited by inter-cell interference, uniform reduction of transmit power has little effect on the NMSE. We skip the numerical results here, while noting the added benefit of energy efficiency with pilot OLPC.

We investigate the impact of BS array size on the linear average of NMSE for all terminals in Fig. 5.4. We observe that the performance of the LS estimator is independent of M, and average NMSE is significantly high (≈ 16 dB) in the absence of pilot OLPC. This is because of the disproportionately large NMSE for cell-edge terminals that dominates the linear average. We observe that the performance of MMSE estimator improves with M, by exploiting additional spatial information contained in larger covariance matrices (red curve).
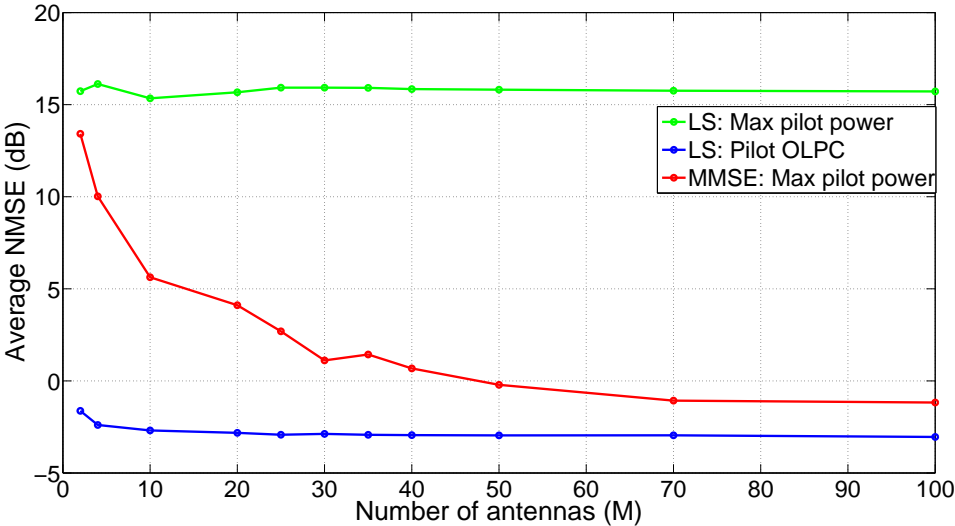

Figure 5.4: Linear average of channel estimation NMSE for all terminals. MMSE estimation improves with M by effectively rejecting inter-cell pilot interference ($L = 21, U = 1, K = 12$)

Forward-link Precoding with Full Pilot Reuse

The channel estimates obtained above are used for precoding at the BS during forward-link transmission. We assume that transmissions from all BSs are synchronized, which is the worst-case scenario for forward-link SINR at the terminals. In Sec. 2.2 we derived the weight vectors for two of the most common linear precoders: MRT and ZF. Also, we obtained the expression for forward-link SINR at the terminal in Eq. 2.12. Here, we evaluate the SINR for our setup, and analyze the impact of contaminated channel estimates on these precoding techniques. We investigate the impact of contaminated channel estimates on the forward-link SINR at the terminal when the $M = 100$-antenna BSs uses linear precoding in each of the $L = 21$ cells. We employ full pilot reuse, pilot OLPC and LS channel estimation.

In Fig. 5.5a, we compare the SINR for contaminated channel estimates (dashed curves) with the knowledge of perfect CSI at the BS (solid curves) for different values of K. In Sec. 2.2 we observed that with MRT precoding, the received signal power attenuates due to *leakage* towards terminals reusing the same pilot sequence in other cells. Moreover, this leakage adds to the inter-cell interference at each terminal. In Fig. 5.5a, the degradation in SINR due to pilot contamination is the horizontal distance between dashed and solid curves of the same color. We observe that for $K = \{3, 6\}$ (black and blue respectively), pilot contamination causes a significant SINR degradation. As K increases to $\{12, 24\}$ (red and green respectively), the impact of contamination on SINR becomes less pronounced. This is because for small K, a large BS array effectively mitigates the intra-cell interference. Therefore, the SINR is dominated by inter-cell interference which is directly affected by pilot contamination. However for larger values of K, the SINR is dominated by intra-cell interference, and the effect of pilot contamination is less.

The effect of contaminated channel estimates on ZF precoding is shown in Fig. 5.5b. Here again, the effect of pilot contamination is large when K is small, which can be explained using the analysis of ZF precoding in Sec. 2.2. The ZF weight vector causes the BS to place nulls in direction of interfering terminals, and transmit the signal in all other directions. With pilot contamination, the channel estimates contain spatial information of the terminals reusing a pilot sequence in other cells. Therefore, the BS is implicitly forced to place nulls in direction of these terminals. This means that the directions in which BS transmits are reduced, attenuating the received signal power at desired terminal. With perfect CSI, nulls are placed in $K - 1$ directions to avoid intra-cell interference. The pilot contamination in case of small K introduces several 'new' directions where the BS places nulls, having a large impact on the SINR.

(a) MRT precoding

(b) ZF precoding

Figure 5.5: Impact of pilot contamination on forward-link terminal SINR with pilot OLPC, LS estimation and linear precoding ($L = 21, U = 1, M = 100$)

We evaluate the forward-link per-cell sum rate $\bar{\mathcal{R}}$ in Fig. 5.6 using the SINR values obtained above. As described in Sec. 5.1, we assume coherence interval of a single LTE RB pair ($14 \times 12 = 168$ orthogonal resources). With full pilot reuse, the number of resources reserved for pilot transmission is equal to K, and the rest are available for data. In case of perfect CSI, we assume an equal amount of pilot overhead for fair comparison with pilot contamination case. In Fig. 5.6a, we

plot $\bar{\mathcal{R}}$ with MRT precoder described in Eq. 2.9 for different values of K and M. We observe that pilot contamination causes more detriment in $\bar{\mathcal{R}}$ for larger M. However even in the presence of pilot contamination, increasing the array size is beneficial as it reduces the intra-cell interference.
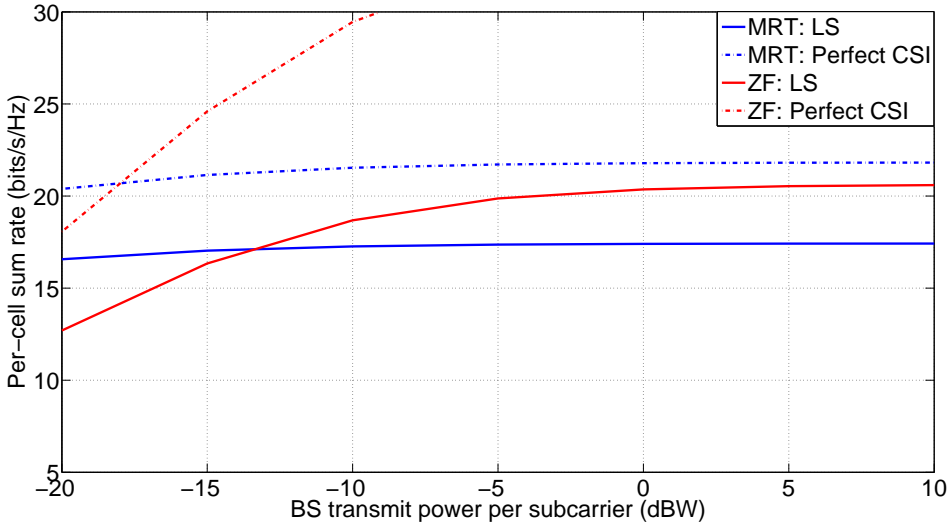

(a) MRT Precoding

(b) ZF Precoding
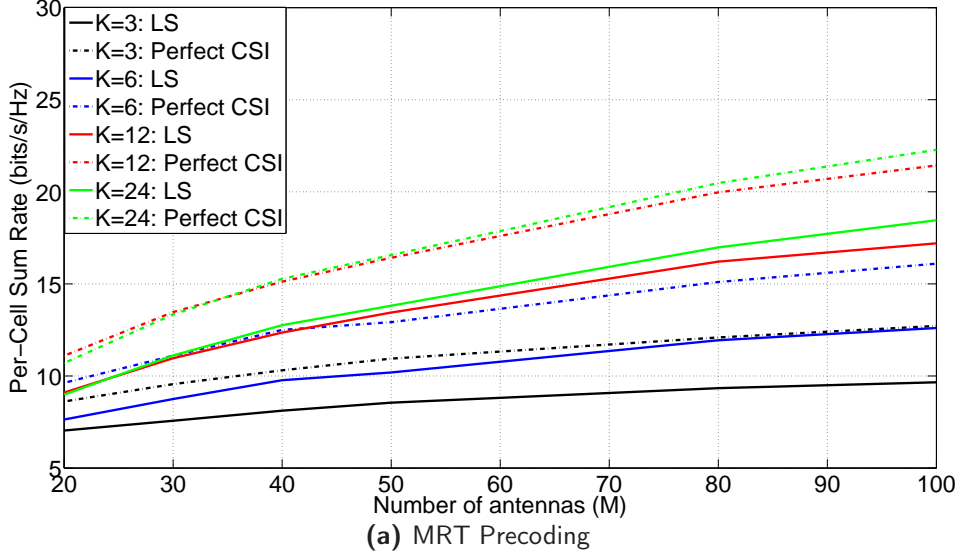
Figure 5.6: Effect of pilot contamination on forward-link sum rate ($L = 21, U = 1$). The sum rate with LS and pilot OLPC (solid curves) is significantly lower than without contamination (dashed curves) for both MRT and ZF. BS transmit power per subcarrier is fixed at $P_{\mathrm{B}} = 0.067$W for all M and K

In Fig. 5.6b, we evaluate $\bar{\mathcal{R}}$ with same parameters as above, but now implementing a ZF precoder at the BS. A ZF precoder attempts to exploit the excess degrees of freedom at the BS to cancel intra-cell interference. Therefore, it performs better than MRT if there are 'large enough' spatial degrees of freedom at the BS. With the ZF precoder described in Eq. 2.9, we observe significant improvement in $\mathcal{R}$ when the array size is much larger than the number of terminals served ($M > 70$). However, for a relatively small array serving a large number of terminals, it leads to a loss in $\bar{\mathcal{R}}$ as compared to the MRT precoder ($M < 30, K = \{12, 24\}$).

From these figures, we observe that the worst-case pilot contamination considered here ($U = 1$, time-synchronized BS) can be a significant impairment to massive MIMO performance. The contamination-free achievable rate is less for MRT precoding as compared to ZF. However, pilot contamination causes a larger impact on ZF than MRT precoding, which results in comparable $\bar{\mathcal{R}}$ with both techniques. The addition of more terminals allows more spatially multiplexed data streams, but also increases the intra-cell interference. We observe that in terms of $\bar{\mathcal{R}}$, increasing K gives progressively diminishing gains.

It is worth reiterating here that ZF performs better than MRT in high SNR (interference-limited) regime. However, it has higher complexity than MRT since it involves matrix inversion. In the presence of pilot contamination, simple MRT processing in a massive MIMO system can effectively mitigate interference, and approach ZF performance even at high BS transmit powers. The effect of different BS transmit power levels on sum rate is shown in Fig. 5.7.

Figure 5.7: Effect of BS transmit power on the performance of MRT and ZF precoding. MRT performs better than ZF in the low-SNR regime. With higher transmit power and pilot contamination, ZF slightly outperforms MRT ($L = 21, U = 1, K = 12, M = 100$)

Earlier, we investigated the MMSE estimator in terms of channel estimation NMSE and found that it provides good channel estimates (NMSE<0dB) for majority of the terminals even in the absence of pilot OLPC. Here, we evaluate $\bar{\mathcal{R}}$ with $K = 12$ and full pilot reuse, when MMSE channel estimates are used for linear precoding. The $\bar{\mathcal{R}}$ with MRT precoding is depicted with blue curves in Fig. 5.8. We observe that MMSE almost completely eliminates pilot contamination, and the slight difference in $\bar{\mathcal{R}}$ as compared to perfect CSI could be because of noise during pilot reception. On the other hand with ZF precoding (red curves), MMSE actually leads to a degradation in $\bar{\mathcal{R}}$ as compared to LS estimates.
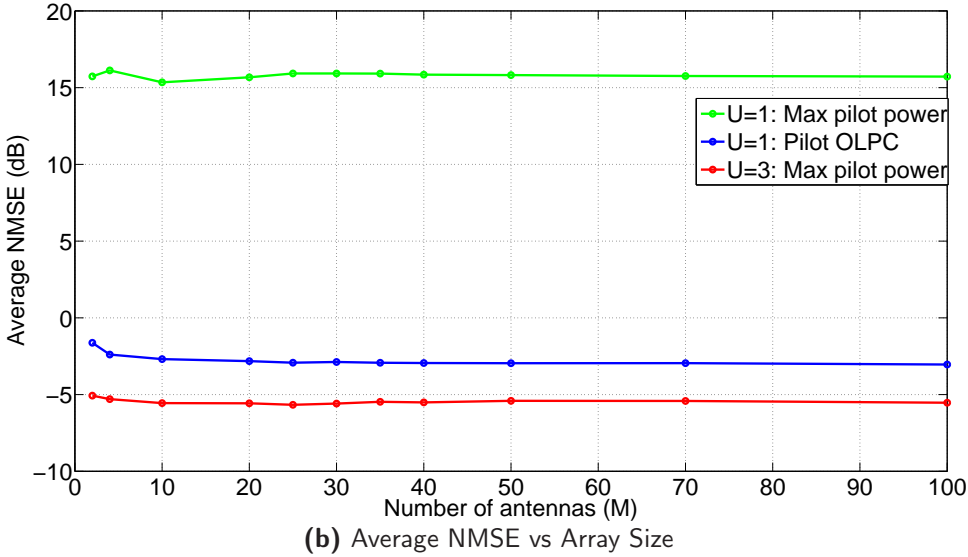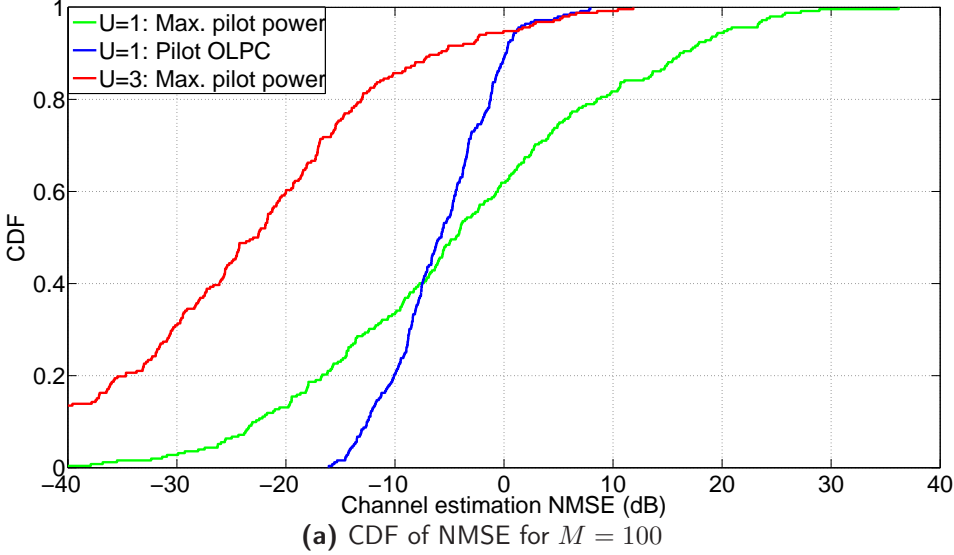

Figure 5.8: Effect of channel estimation technique on sun rate. MMSE performs better than LS with MRT , but under-performs LS with ZF. We use pilot OLPC with LS to keep channel estimation error below 0 dB. ($L = 21, U = 1$)

Less Aggressive Pilot Reuse

An intuitive approach to reduce pilot contamination and thus improve channel estimates is to incorporate a less aggressive pilot reuse scheme as discussed in Sec. 4.2. The cost of this approach is a higher overhead of pilot resources which limits the coherent resources available for forward- and reverse-link data transmission. In Fig. 5.9a, we evaluate the distribution of channel estimation NMSE for LS estimator with $U = 3$ and $K = 12$ (red curve). We compare this against the full pilot reuse scenario discussed above, with and without pilot OLPC (blue and green curves respectively).

We observe that even without power control, the NMSE significantly improves with $U = 3$. The NMSE for most terminals falls below 0 dB, implying they have good enough channel estimates to be used for forward-link precoding and reverse-link processing at the BS. We observe that even for $U = 3$, there is significant variation in the NMSE across terminals in the absence pilot OLPC. We observe

however, that this variation is the logarithmic scale with all values smaller than 0dB. In linear scale, the corresponding values of NMSE are quite small, and the absolute variation is much smaller compared to $U = 1$. Thus with a less aggressive pilot reuse, we might not need pilot OLPC at all, avoiding the associated overhead described in Sec. 4.1.

(a) CDF of NMSE for $M = 100$

(b) Average NMSE vs Array Size

Figure 5.9: Effect of less aggressive pilot reuse on NMSE of LS channel estimation. The pilot contamination is effectively mitigated, even in the absence of pilot OLPC ($L = 21, K = 12$)

In Fig. 5.10a and 5.10b, we investigate the forward-link SINR at the terminals with $U = 3$ and LS channel estimation for $K = 12$. We observe that in case of both MRT and ZF precoding, $U = 3$ almost completely removes the impact of pilot contamination. This implies that the pilot contamination from adjacent cells constitutes almost all of the SINR degradation in our massive MIMO setup.

(a) MRT precoding
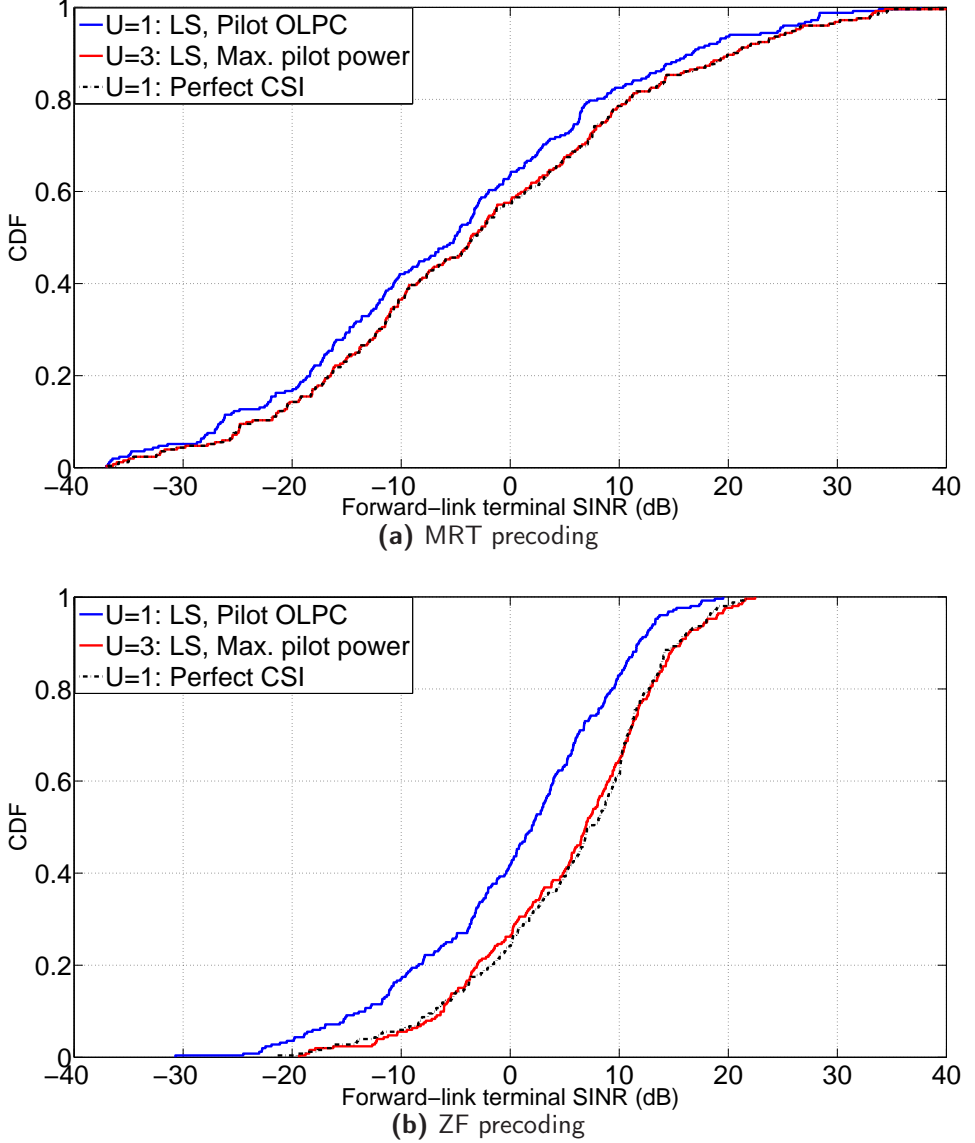
(b) ZF precoding

Figure 5.10: Effect of less aggressive pilot reuse on forward-link SINR at the terminal. $U = 3$ almost completely eliminates the impact of pilot contamination on MRT/ZF precoding ($M = 100, L = 21, K = 12$)

Clearly, the gains in SINR with $U = 3$ are at cost of additional pilot over-head, and we need to account for this while evaluating the system performance. We investigate the pilot-data symbol trade-off in Fig. 5.11a and 5.11b, where we illustrate the achievable rate per terminal with $L = 21$ and $K = 12$.

(a) MRT precoding

(b) ZF precoding
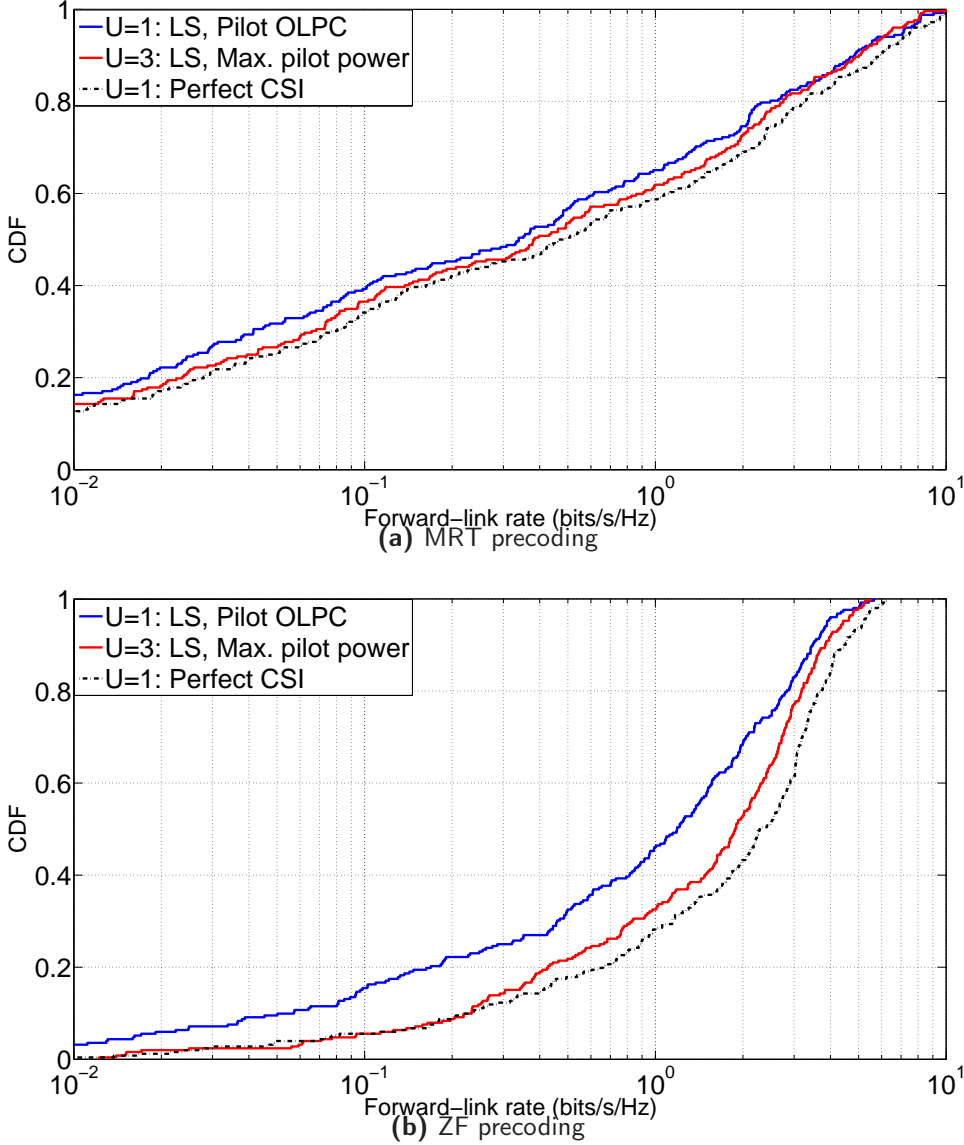
Figure 5.11: Effect of less aggressive pilot reuse on forward-link rate per terminal. The rate increases for our assumed coherence interval of 1 LTE RB pair ($L = 21, K = 12, M = 100$)

With $U = 3$, a total of $3K$ resources are reserved for pilot transmission, and the rest are available for data. The first set of K pilot resources is allocated to terminals within cells $\{1, 4, 7, ...\}$, the second set to cells $\{2, 5, 8, ...\}$, etc. We observe that for our assumed coherence interval of 1 LTE RB pair, the rate achieved by each terminal increases. In other words, the gains of improved channel estimation with $U = 3$ outweigh the additional pilot overhead. If the coherence interval is larger, the fraction of resources wasted on pilots is reduced, leading to higher rates. However for shorter coherence intervals or larger number of terminals, the pilot-data trade-off may become unfavorable, and lead to a drop in performance.
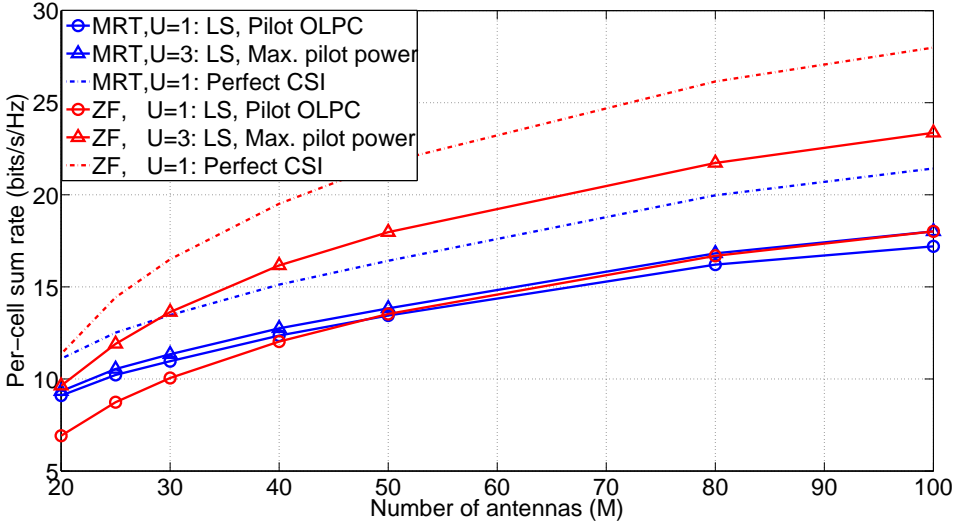
Figure 5.12: Per-cell forward-link sum rate for pilot reuse 1/3 with MRT/ZF precoding $(L = 21, K = 12)$

In Fig. 5.12, we compare $\bar{\mathcal{R}}$ against different BS array sizes for $U = 1$ and 3, and observe that the it benefits from less aggressive pilot reuse for all M.

Next, we evaluate the soft pilot reuse (SPR) scheme described in Sec. 4.2 in the current setup. We implement a 1/2 SPR scheme that groups the terminals based on the knowledge of their reverse link SNR while assigning pilots. The terminals with strong channels to the BS (cell-center terminals) use same set of pilots in each cell, $U_{\mathrm{C}} = 1$. The cell-edge terminals are expected to be at the risk of higher pilot interference. Therefore, the terminals with lowest half of the SNR are assumed to be at the cell edge and assigned a less aggressive pilot reuse factor of $U_{\mathrm{E}} = 3$. The total number of resources used for pilot transmission are therefore $0.5KU_{\mathrm{C}} + 0.5KU_{\mathrm{E}} = 2K$. We evaluate the forward-link SINR at the terminals with ZF precoding and 1/2 SPR in Fig. 5.13a. We observe that this scheme improves the SINR as compared to full pilot reuse at the cost of K additional resources. Also, it underperforms $U = 3$ scheme, that uses an additional K resources.

(a) SINR, ZF precoding

(b) Achievable per-cell sum rate, ZF precoding
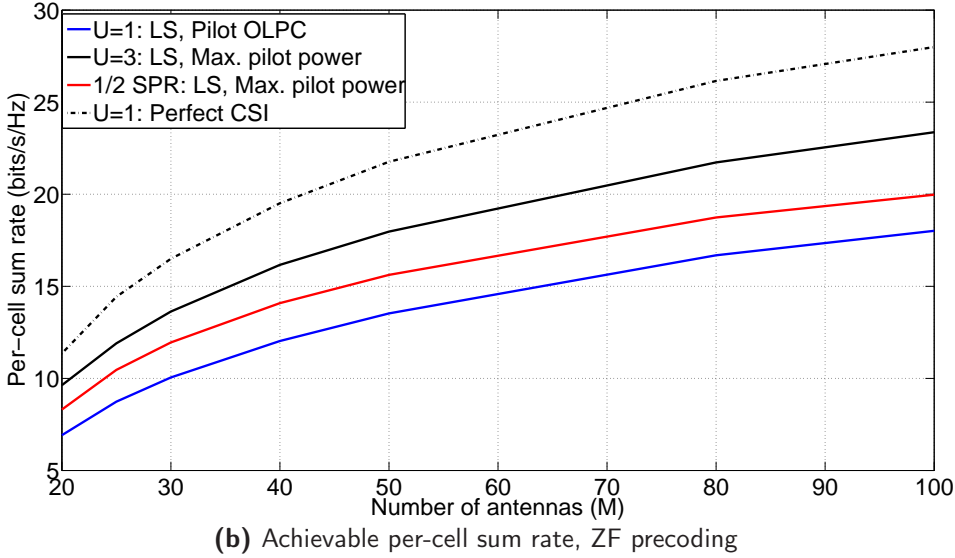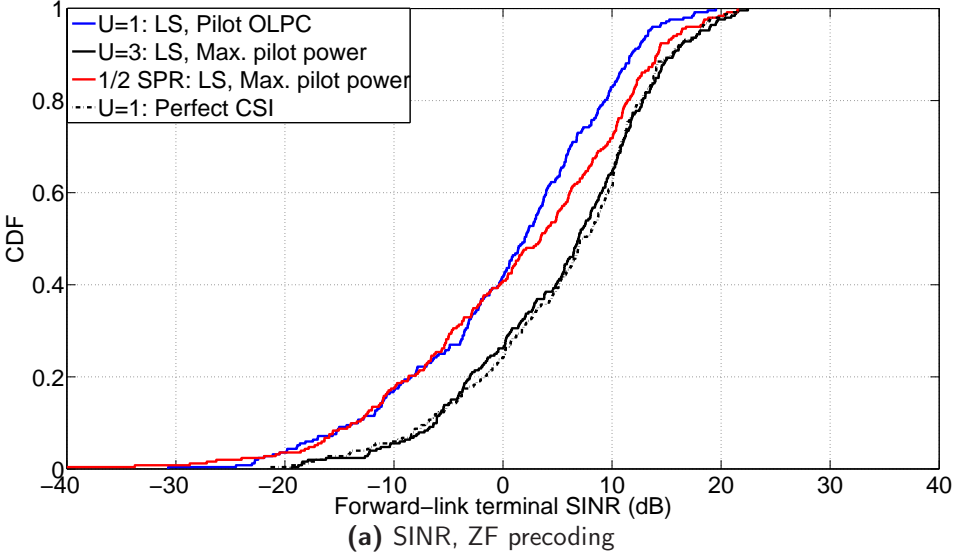
Figure 5.13: Forward-link SINR and rate for soft pilot reuse (SPR) with LS channel estimation ($L = 21, K = 12, M = 100$)

To account for the difference in pilot overhead, we investigate the achievable rate for each technique. In Fig. 5.13b, we compare the per-cell sum rate obtained with 1/2 SPR against $U = 1$ and $U = 3$ with LS channel estimation and ZF precoding. We observe that for the assumed coherence interval of 2 LTE RBs, 1/2 SPR improves the rate over full pilot reuse. This means that the reduction in pilot contamination with 1/2 SPR trades off favorable with the increased pilot

overhead. However, $U = 3$ provides additional improvement in rates over $1/2$ SPR for all M. It must be pointed out however, that the rate depends on the size of coherence interval and the number of terminals to be served. For a shorter coherence interval or a larger number of terminals, the additional pilot overhead with $U = 3$ might no longer remain favorable in terms of pilot-data trade-off.
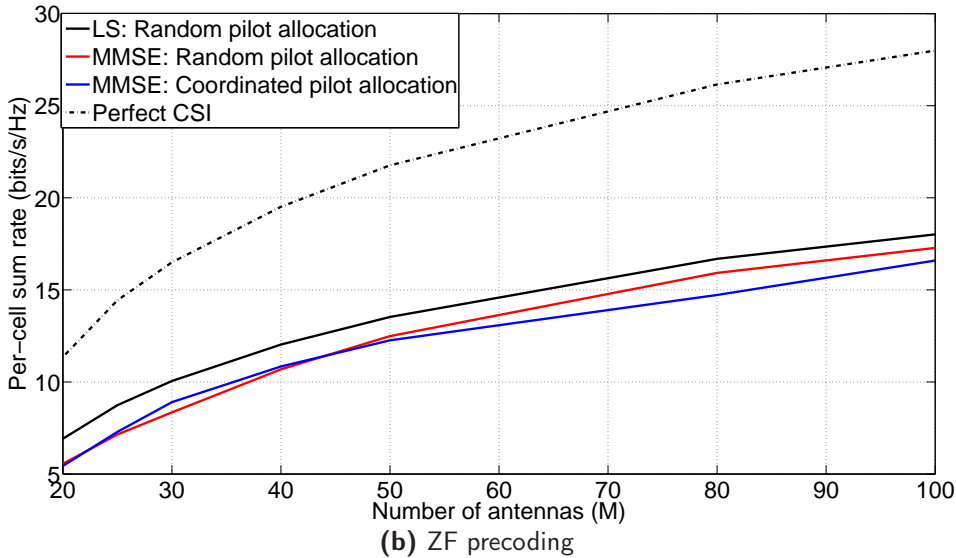
Coordinated Pilot Allocation

The discussion till now has utilized random allocation of pilots to terminals within each cell. With this, we have analyzed the pilot contamination effect and determined the conditions under which it can be mitigated using simple power control and pilot reuse schemes. As briefly noted in Fig. 5.3, an MMSE estimator can be used to obtain improved channel estimates by exploiting spatial information contained in the cross-channel covariance matrices. Additionally, the covariance matrices can be exploited for improved pilot allocation as discussed in Sec. 4.3.
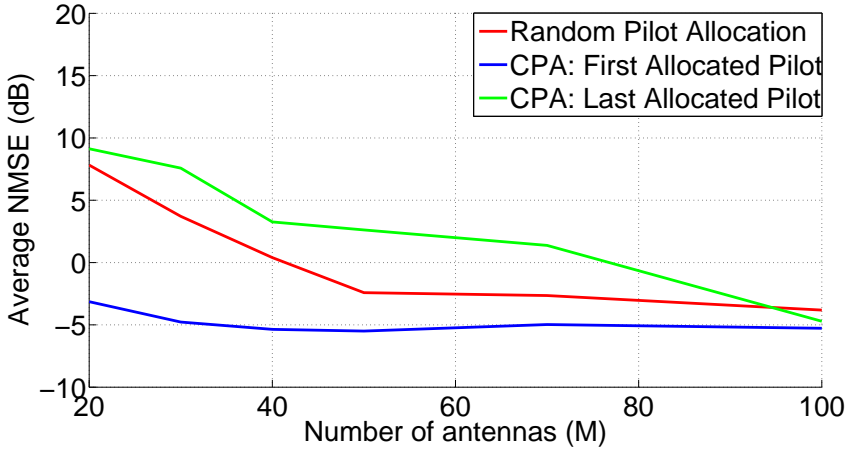
Figure 5.14: Avg. channel estimation NMSE for terminals reusing a pilot sequence. CPA provides gains over random allocation only for the first few pilots. $(L = 21, U = 1, K = 12, M = 100)$

We investigate the efficacy of this scheme in Fig. 5.14 for the $L = 21$-cell setup, when the same set of $K = 12$ orthogonal pilot sequences are distributed among the terminals within each cell. We plot the average NMSE for all terminals allocated the first pilot sequence with a blue curve. We observe that in this case, the average NMSE is significantly lower than that for random pilot allocation (red curve). This is because for the first pilot allocation, the CPA algorithm scans all terminals and picks the most spatially well-separated terminals that reuse this pilot sequence. Thus even for smaller array sizes, CPA is able to effectively eliminate pilot contamination. However, we observe that the average NMSE for last pilot allocation with CPA (green curve) is actually worse than random allocation. This is because of the 'greedy' allocation approach that picks the spatially best-separated terminals in each iteration, leaving poorer choices for subsequent

pilot allocations. Consequently, the last few pilots are forced to be allocated to terminals that overlap significantly in the angular domain. We observe that for larger M, random pilot allocation performs nearly as well as CPA as the channels between a BS and the terminals become more orthogonal.

(a) MRT precoding

(b) ZF precoding

Figure 5.15: Per-cell forward link sum rate using MMSE estimator and coordinated pilot allocation (CPA) ($L = 21, K = 12$)

In Fig. 5.15a and 5.15b, we evaluate the achievable rate CPA with MMSE

channel estimation. We observe that for both MRT and ZF precoding, the CPA technique does not provide any gains over random allocation in terms of sum rate. For the terminals allocated the first few pilots using CPA, the achievable rate certainly increases. However, these gains are offset by the poor channel estimates and thus lower rate for terminals assigned the later pilot sequences. Therefore on average for all pilot allocation in a cell, there is no improvement in sum rate by using CPA.

Conclusions and Further Work

In this thesis, we investigated the impact of pilot contamination on reverse-link channel training and forward-link transmission in $L = 21$ synchronous cells with up to $M = 100$ antennas at the BS and up to $K = 24$ terminals per cell. We used the LTE forward-link grid while assuming a coherence interval of 14 OFDM symbols (1 ms) and 12 subcarriers (180 kHz). We evaluated the BS array size and number of terminals for which pilot contamination may cause a significant impairment in system performance. We developed low-complexity methods to mitigate this effect, and compared it with a (high-complexity) inter-cell coordination technique. The key findings of this thesis are summarized below:

1. The performance of LS channel estimator is independent of M, and average normalized mean square error (normalized MSE, or NMSE) of channel estimation is significantly high ($\approx$ 16 dB) in the absence of pilot power control.

2. With pilot open loop power control (pilot OLPC) the NMSE-fairness is improved, and the channel estimates for most terminals are good enough for precoding (NMSE<0 dB).

3. The channel estimation performance of minimum MSE (MMSE) estimator in terms of average NMSE improves with larger array size M, by exploiting additional spatial information contained in the larger $M \times M$ channel covariance matrices.

4. For $K = \{3, 6\}$ and maximum ratio transmission (MRT), pilot contamination causes a significant impairment in forward-link terminal SINR. As K increases to $\{12, 24\}$, the impact of contamination on SINR becomes less pronounced. This is because for small K, a large BS array effectively mitigates the intra-cell interference. Therefore, the SINR is dominated by inter-cell interference which is directly affected by pilot contamination. However for larger values of K, the SINR is dominated by intra-cell interference, and the effect of pilot contamination is less.

5. Similar to MRT, in the case of zero-forcing (ZF) precoding the effect of pilot contamination is large when K is relatively small

6. Even in the presence of pilot contamination, increasing the array size is beneficial in terms of per-cell forward link sum rate $\bar{\mathcal{R}}$ for both MRT and

ZF precoding. However, the worst-case pilot contamination (pilot reuse $U = 1$) degrades $\bar{\mathcal{R}}$ significantly compared to contamination-free channel estimates for larger array sizes.

7. With perfect channel estimates, the performance of ZF increases dramatically with the BS transmit power, while the performance of MRT remains almost constant. However with pilot contamination, the $\bar{\mathcal{R}}$ with ZF saturates quickly with high-enough BS transmit power.

8. The channel estimation NMSE improves dramatically with less aggressive pilot reuse ($U = 3$), even in the absence of any pilot power control. The NMSE for most terminals falls below 0dB, implying they have good enough channel estimates to be used for forward-link precoding and reverse-link processing at the BS.

9. In case of both MRT and ZF precoding, $U = 3$ almost completely removes the impact of pilot contamination on forward-link terminal SINR, at the cost of additional resources spend on pilot transmission.

10. For our setup, $\bar{\mathcal{R}}$ increases in case of $U = 3$ compared to $U = 1$. In other words, the gains of improved channel estimation with $U = 3$ outweigh the additional pilot overhead.

11. The 1/2 SPR scheme improves the forward-link terminal SINR as compared to full pilot reuse, at the cost of K additional pilot resources.

12. The 1/2 SPR scheme improves $\bar{\mathcal{R}}$ as compared to $U = 1$. This means that the reduction in pilot contamination with 1/2 SPR trades off favorable with the increased pilot overhead. However, $U = 3$ provides additional improvement in rates over 1/2 SPR for all M.

13. For the first few pilots allocations, inter-cell coordinated pilot allocation (CPA) gives large NMSE gains over random pilot allocation. However for the last few pilots the average NMSE is severely degraded in case of CPA. This is because CPA works on a 'greedy' approach that picks the spatially best-separated terminals for first pilot allocation, leaving fewer choices for rest of the pilots.

14. For both MRT and ZF precoding, the CPA technique does not provide any gains over random allocation in terms of $\bar{\mathcal{R}}$.

Further Work

The study of massive MIMO systems is still in its initial stages, and there are several interesting techniques that may be used to improve the system performance/ make it feasible to deploy in practice. We enumerate a few areas that are worth exploring in the context of the work completed during this thesis.

1. In our discussion, we have assumed perfect orthogonality of the pilot sequences, which might not be true in practice. Additionally, it is interesting to explore the case of non-orthogonal pilot sequences from the perspective of fewer number of resources spent on channel estimation.

2. We have identified that for forward-link transmission, the impact of pilot contamination on rate is most pronounced for small number of served terminals. It is interesting to determine the values of K and M under which it is beneficial to adopt various mitigation techniques.

3. The current thesis focuses on the impact of pilot contamination on forward-link transmission. It might be interesting to explore the reverse-link processing of data signals at the BS as well, which is expected to have different SINR characteristics.

4. The mitigation techniques discusses here can be used to formulate improved receivers in the reverse link, that can be used for better subsequent channel estimates.

5. We have considered simple MRT and ZF precoders, that are extreme approaches in the sense of maximizing the signal power or minimizing the intra-cell interference. A mid-way scheme like regularized zero-forcing can be explored in the context of massive MIMO as well.

6. CPA gives good channel estimation gains only for the first few pilots on account of the greedy approach. It might be interesting to explore better algorithms for pilot allocation, that give a result closer to the global optimum.

References

[1] A. Paulraj, D. Gore, R. Nabar, and H. Bolcskei, "An overview of MIMO communications - a key to Gigabit wireless," *Proc. IEEE*, vol. 92, no. 2, pp. 198–218, Feb. 2004.

[2] T. L. Marzetta, "How much training is required for multiuser MIMO?" in *Proc. IEEE Asilomar Conf. Signals, Systems, and Computers*, Oct 2006, pp. 369–363.

[3] B. Gopalakrishnan and N. Jindal, "How much training is required for multiuser MIMO?" in *Proc. IEEE Workshop on Signal Processing Adv. in Wireless Commun. (SPAWC)*, Jun 2011, pp. 381–385.

[4] N. Krishnan, R. Yates, , and N. Mandayam, "Uplink linear receivers for multicell multiuser MIMO with pilot contamination: Large system analysis," *IEEE Trans. Wireless Commun.*, vol. 13, no. 8, pp. 4360–4373, Aug. 2014.

[5] A. Paulraj, R. Nabar, and D. Gore, *Introduction to Space-Time Wireless Communications*. Cambridge University Press, 2008.

[6] E. Dahlman, S. Parkvall, and J. Sköld, *4G: LTE/LTE-Advanced for Mobile Broadband*. Elsvier, 2014.

[7] ITU-R, "Guidelines for evaluation of radio interface technologies for IMT-Advanced," International Telecommunication Union, Tech. Rep. ITU-R, M.2135-1, Dec. 2009.

[8] M. Steinbauer, A. Molisch, and E. Bonek, "The double-directional radio channel," *IEEE Anten. and Prop. Mag.*, vol. 43, no. 4, pp. 51–63, Aug. 2001.

[9] D. Chizhik, F. Rashid-Farrokhi, J. Ling, and A. Lozano, "Effect of antenna separation on the capacity of BLAST in correlated channels," *IEEE Commun. Lett.*, vol. 4, no. 11, pp. 337–339, Nov. 2000.

[10] A. Scherb and K. Kammeyer, "Bayesian channel estimation for doubly correlated mimo systems," in *Proc. ITG Workshop on Smart Antennas (WSA)*, Vienna, Austria, Feb 2007.

[11] M. Frenne, "On RSRP determination for UE to node association," 3rd Generation Partnership Project, Tech. Rep. 3GPP, TSG-RAN WG1 Meeting 75, Aug. 2001.

[12] S. Järmyr, "MIMO transceiver design for multi-antenna communications over fading channels," Ph.D. dissertation, Royal Institute of Technology (KTH), Sweden, 2013.

[13] J. Duplicy, B. Badic, R. Balraj, R. Ghaffar, P. Horvath, F. Kaltenberger, R. Knopp, I. Kovacs, H. Nguyen, D. Tandur, and G. Vivier, "MU-MIMO in LTE systems," *EURASIP J. on Wireless Comm. and Netw.*, Mar. 2011.

[14] T. Marzetta, "Noncooperative cellular wireless with unlimited numbers of base station antennas," *IEEE Trans. Wireless Commun.*, vol. 9, no. 11, pp. 3590–3600, Sep. 2010.

[15] E. Larsson, O.Edfors, F.Tufvesson, and T. Marzetta, "Massive MIMO for next generation wireless systems," *IEEE Commun. Mag.*, vol. 52, no. 2, pp. 186–195, Feb. 2014.

[16] G. Xiang, F. Tufvesson, O. Edfors, and F. Rusek, "Measured propagation characteristics for very-large MIMO at 2.6 GHz," in *Proc. IEEE Asilomar Conf. Signals, Systems, and Computers*, Nov. 2012, pp. 295–299.

[17] M. Biguesh and A. Gershman, "Training-based mimo channel estimation: a study of estimator tradeoffs and optimal training signals," *IEEE Trans. Signal Process.*, vol. 54, no. 3, pp. 884–893, Mar 2006.

[18] M. Jain, J. Choi, T. Kim, D. Bharadia, S. Seth, K. Srinivasan, P. Levis, S. Katti, and P. Sinha, "Practical, real-time, full duplex wireless," in *Proc. Int. Conf. on Mobile Computing and Networking (ACM)*, New York, NY, USA, 2011, pp. 301–312.

[19] F. Kaltenberger, J. Haiyong, M. Guillaud, and R. Knopp, "Relative channel reciprocity calibration in mimo/tdd systems," in *Proc. Future Network and Mobile Summit*, Jun 2010, pp. 1–10.

[20] G. Raleigh, S. Diggavi, V. Jones, and A. Paulraj, "A blind adaptive transmit antenna algorithm for wireless communication," in *Proc. IEEE Int. Conf. on Commun. (ICC)*, vol. 3, Jun 1995, pp. 1494–1499.

[21] B. Hochwald and T. Maretta, "Adapting a downlink array from uplink measurements," *IEEE Trans. Signal Process.*, vol. 49, no. 3, pp. 642–653, Mar 2001.

[22] S. M. Kay, *Fundamentals of statistical signal processing: estimation theory.* Prentice-Hall, 1993.

[23] H. Yin, D. Gesbert, M. Filippou, and Y. Liu, "A coordinated approach to channel estimation in large-scale multiple-antenna systems," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 2, pp. 264–273, Feb 2013.

[24] J. Jose, A. Ashikhmin, T. Marzetta, and S. Vishwanath, "Pilot contamination problem in multi-cell TDD systems," in *Proc. IEEE Int. Symp. on Inf. Theory (ISIT)*, Jun 2009, pp. 2184–2188.

[25] P. Xu, J. Wang, and J. Wang, "Effect of pilot contamination on channel estimation in massive MIMO systems," in *Proc. Int. Conf. on Wireless Commun. Signal Proces. (ICWCSP)*, Oct 2013, pp. 1–6.

[26] A. Simonsson and A. Furuskar, "Uplink power control in LTE - overview and performance," in *Proc. IEEE Veh. Tech. Conf. (VTC)*, Sep 2008, pp. 1–5.

[27] S. Sesia, I. Toufik, and M. Baker, *LTE – The UMTS Long Term Evolution: From Theory to Practice.* John Wiley & Sons, 2011.

[28] D. Huo, "Clarification on the wrap-around hexagon network structure," www.ieee802.org/20/Contribs/C802.20-05-15.doc, Mar 2005.

[29] F. Rusek, D. Persson, B. K. Lau, E. Larsson, T. Marzetta, O. Edfors, and F. Tufvesson, "Scaling up MIMO: Opportunities and challenges with very large arrays," *IEEE Signal Process. Mag.*, vol. 30, no. 1, pp. 40–60, Jan. 2013.

[30] 3GPP, "Spatial channel model for multiple input multiple output (mimo) simulations," 3GPP, Tech. Rep. 3GPP, TR25.996 V12.0.0, Sep. 2014.

Appendix

7.1 MMSE channel estimator

We assume that a pilot sequence $\mathbf{s}_k$ of length τ symbols is assigned to the k^{th} terminal in each of the L cells. The received signal at j^{th} BS due to transmission of this sequence is given by Eq. 3.2, and can be expanded to

$$\mathbf{X}_{jk} = \begin{bmatrix} \mathbf{g}_{jk1} & \mathbf{g}_{jk2} & \cdots & \mathbf{g}_{jkL} \end{bmatrix} \begin{bmatrix} s_{k1} \\ s_{k2} \\ \vdots \\ s_{k\tau} \end{bmatrix} + \mathbf{N}_j. \tag{7.1}$$

We stack the received signal, the channel and the noise into column vectors $\mathbf{x}_{jk}$, $\mathbf{g}_{jk}$ and $\mathbf{n}_j$ respectively. The equation can then be represented as

$$\mathbf{x}_{jk} = \mathbf{S}_k \mathbf{g}_{jk} + \mathbf{n}_j, \tag{7.2}$$

where $\mathbf{S}_k$ is the pilot matrix obtained by taking Kronecker product of the pilot sequence with an identity matrix of size M, i.e.

$$\mathbf{S}_k = \begin{bmatrix} s_{k1} \otimes \mathbf{I}_M & s_{k2} \otimes \mathbf{I}_M & \ldots & s_{k\tau} \otimes \mathbf{I}_M \end{bmatrix}. \tag{7.3}$$

The MMSE estimator for channel in the j^{th} cell is then given by [22][23]

$$\widehat{\mathbf{g}}_{jk}^{MMSE} = \mathbf{R}_{jk} \mathbf{S}_k^H \left(\mathbf{S}_k \left(\sum_{l=1}^{L} \mathbf{R}_{jkl} \right) \mathbf{S}_k^H + \sigma_n^2 \mathbf{I}_{\tau M} \right)^{-1} \mathbf{x}_{jk}. \tag{7.4}$$

Applying the matrix inversion identity $\mathbf{A}(\mathbf{I} + \mathbf{BS})^{-1} = (\mathbf{I} + \mathbf{AB})^{-1}\mathbf{A}$, we obtain a simplified expression for the MMSE estimator, we get

$$\widehat{\mathbf{g}}_{jk}^{MMSE} = \mathbf{R}_{jk} \left(\sigma_n^2 \mathbf{I}_M + \tau \sum_{l \in \mathcal{J}} \mathbf{R}_{jkl} \right)^{-1} \mathbf{S}_k^H \mathbf{x}_{jk}. \tag{7.5}$$