



Terms and Conditions of Use of Digitised Theses from Trinity College Library Dublin

Copyright statement

All material supplied by Trinity College Library is protected by copyright (under the Copyright and Related Rights Act, 2000 as amended) and other relevant Intellectual Property Rights. By accessing and using a Digitised Thesis from Trinity College Library you acknowledge that all Intellectual Property Rights in any Works supplied are the sole and exclusive property of the copyright and/or other IPR holder. Specific copyright holders may not be explicitly identified. Use of materials from other sources within a thesis should not be construed as a claim over them.

A non-exclusive, non-transferable licence is hereby granted to those using or reproducing, in whole or in part, the material for valid purposes, providing the copyright owners are acknowledged using the normal conventions. Where specific permission to use material is required, this is identified and such permission must be sought from the copyright holder or agency cited.

Liability statement

By using a Digitised Thesis, I accept that Trinity College Dublin bears no legal responsibility for the accuracy, legality or comprehensiveness of materials contained within the thesis, and that Trinity College Dublin accepts no liability for indirect, consequential, or incidental, damages or losses arising from use of the thesis for whatever reason. Information located in a thesis may be subject to specific use constraints, details of which may not be explicitly described. It is the responsibility of potential and actual users to be aware of such constraints and to abide by them. By making use of material from a digitised thesis, you accept these copyright and disclaimer provisions. Where it is brought to the attention of Trinity College Library that there may be a breach of copyright or other restraint, it is the policy to withdraw or take down access to a thesis while the issue is being resolved.

Access Agreement

By using a Digitised Thesis from Trinity College Library you are bound by the following Terms & Conditions. Please read them carefully.

I have read and I understand the following statement: All material supplied via a Digitised Thesis from Trinity College Library is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of a thesis is not permitted, except that material may be duplicated by you for your research use or for educational purposes in electronic or print form providing the copyright owners are acknowledged using the normal conventions. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone. This copy has been supplied on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.



Trinity College Dublin
Coláiste na Tríonóide, Baile Átha Cliath
The University of Dublin


PHD THESIS

**Enhanced Multi-Antenna Capabilities
through Small Cell Collaboration**

Author:
Danny FINN

Supervisor:
Prof. Luiz A. DASILVA

5th May 2015




Thesis 20675

Declaration

I declare that this thesis has not been submitted as an exercise for a degree at this or any other university and is entirely my own work.

I agree to deposit this thesis in the University's open access institutional repository or allow the Library to do so on my behalf, subject to Irish Copyright Legislation and Trinity College Library conditions of use and acknowledgement.

Signed:



Danny Finn

5th May 2015

Summary

This thesis shows that collaboration between base stations in small cell scenarios can lead to enhanced multiple-antenna gains. We demonstrate this for the case of *Multi-User MIMO-based reassignments* between adjacent small cells, from the perspectives of spectral efficiency, energy efficiency and number of spatial layers effectively utilised.

Downlink Multi-User MIMO (MU-MIMO) is a spatial multiplexing technique in which multiple transmit antennas at an evolved Node Base station (eNB) are used to simultaneously serve multiple User Equipments (UEs), each on a different spatial layer, in such a way that the interference between the beams directed at co-scheduled UEs is kept as low as possible. Due to MU-MIMO's ability to provide large spatial multiplexing gains without requiring additional antennas on the UE, and its ability to overcome rank deficiency problems (which often limit single point-to-point spatial multiplexing), MU-MIMO capabilities have been highly emphasised in recent standardisation.

In this thesis we focus on the use of MU-MIMO in coordinated LTE small cell networks. Small cells are an appealing technology to operators due to their low costs of deployment and operation, and their ability to dramatically increase the spatial reusability of limited spectrum.

MU-MIMO performs best when the antennas at the eNB experience highly correlated channel characteristics, which is often the case in small cell scenarios. However, for effective operation, MU-MIMO requires the channels between the base station and different simultaneously-served UEs to be sufficiently uncorrelated (close to orthogonal) to ensure that excessive cross-layer interference is not experienced. While this cross-layer (or multi-user) interference can be suppressed at the transmitter using techniques such as Zero-Forcing Beam Forming, there is an associated cost in terms of signal power, meaning that the combined throughput of the spatial layers of multiple non-orthogonal UEs may be less than the

unsuppressed throughput of a single UE transmitted to on single layer alone. In macrocell networks this problem is overcome by taking advantage of multi-user diversity through MU-MIMO scheduling. On the other hand, in small cell scenarios where the number of UEs per cell is low, sets of UEs with mutually orthogonal channels do not always exist in a given cell and the spatial multiplexing gains of MU-MIMO cannot always be realised. However, due to the high density of deployment in small cell networks, UEs are often in range of multiple small cells at once.

For this reason, we propose MU-MIMO-based reassignments of UEs between adjacent small cells, through which the channel orthogonality between UEs can be increased along with the number of spatial layers which can be effectively utilised. We propose a number of network-wide-mechanisms which target different performance objectives, such as spectral efficiency and energy efficiency, through the use of centralised coordination between neighbouring small cells. We collectively call these mechanisms *Multi-User MIMO across Small Cells*.

This dissertation firstly reviews the background and relevant literature in the areas of MIMO, small cells, multi-cell coordination, and cross-overs between these. We then outline our simulation methodology and lessons learnt, including Extremely Dense Network baseline simulator comparison and selection guidelines, and simulator extensions implemented as part of this work. Finally, we specify, and assess the performance of, our two MU-MIMO-based reassignment mechanisms. Performance is assessed relative to the current practice of assigning UEs to the small cell of highest received power (not taking into account MU-MIMO considerations).

We find that from the reassignment of a single UE, the set of UEs for which MU-MIMO is enabled achieves spectral efficiency average gains of 0.6 (from 2.16 to 2.78) and 0.7 (from 1.43 to 2.20) bps/Hz for indoor and outdoor deployments, corresponding to the enabling of MU-MIMO for the reassigned UE on 96% and 94% of scheduled subbands, respectively. This affected UE set for which MU-MIMO is enabled consists of the reassigned UE and the UE with which they get paired in the target cell. We also found that, through MU-MIMO-based reassignments, we can reduce the required number of active small cells by more than 25% (indoor) and 35% (outdoor), while still achieving increases in spectral efficiency of the affected UE set.

Acknowledgements

Firstly and foremost, huge thanks are due to my supervisor Prof. Luiz DaSilva who always had faith in my abilities when I thought a PhD was too big of an undertaking, who constantly sought rigour and perfection in every aspect of my work, who instantly knew the perfect phrasing of every key sentence, and who always went the extra mile in making sure that CTVR was an ideal place to work.

Secondly, to Dr. Hamed Ahmadi, for working so closely with me, directing me each time I was lost, for being a huge inspiration throughout the longevity of my PhD, and for introducing me to more people than I can remember.

And thirdly, to Dr. Andrea Cattoni, for helping to kick-start my work in the area of MU-MIMO-based small cell reassignments which led to the main contributions of this thesis, for constantly providing constructive feedback, for always expecting others who work with him to (try to) work as hard as him, and for making sure that my time in Aalborg was such a worthwhile and enjoyable experience.

My deepest of gratitude is due to the entire team at CTVR. To Prof. Linda Doyle for being the most enthusiastic professor a department could have. To Pedro and Francisco, for answering so many of my questions before and during the writing of this thesis. *Malandro mesmo é o cavalo marinho, que finge ser peixe pra não ter que puxar carroça.* To Paolo, for being a true Italian. To Quentin, for being a true *French*. To David for being a great student to work with. To Arman, for finding incomprehensible ways of explaining the simplest of things. To the builders at my window, for making sure to work just as loudly when I come in on Sunday as any other day of the week. To Emanuele, Carlo, Jonathan, Nicola, Ahmed, Seamas, Irene, Jacek, Johann, Eamonn, Francesco, Nick, Ioannis, Jasmina, Elma, and Yong for being great people to work and share an office with, you are all an inspiration. To Justin, for working with me across so many years, for pushing me toward the end, and most importantly for making a PhD look easy. And to all the rest of the

ever-growing close-knit team.

Thanks are due to the crews from CREW and Aalborg, as well as Rouzbeh and Holger in Alcatel Lucent, for fruitful collaborations of which I'm sure there'll be more to come.

To Muireann, for making thesis writing a lot less daunting.

To the Dublin University Trampoline Club, to which I will be forever indebted for so hugely enhancing my college experience to the extent that I would almost consider Trinity College and the Trampoline Club synonymous. The club and the people that made it have provided 8 years of great training, great coaching, great trips away, great competitions, great nights out, and great friends, which I will never forget. Special mentions to Pa and Conor O'Brien, for being a great flatmate, and visiting at just the right time, respectively.

To Bevan, Rob, Glenn, Kieron and Will for providing breaks when I needed them.

To my parents, for raising a great family and dedicating so much time to each of us. You have made sure that there is nothing we can't do if we put our minds to it.

To my examiners, Prof. Simon Saunders and Dr. Muhammad Ali Imran, for their in-depth feedback which greatly strengthened the contributions of this work.

Most importantly of all, I dedicate this thesis to Fiona for keeping me focused, bearing with me as the end approached, and always providing the best of company.

Contents

Declaration	i
Summary	iii
Acknowledgements	v
Contents	vii
List of Figures	xi
List of Tables	xiii
1 Introduction	1
1.1 Overview	1
1.2 Key Contributions	4
1.2.1 MU-MIMO-based Reassignment Mechanism aimed toward Spectral Efficiency Gains	4
1.2.2 Refined MU-MIMO-based Reassignment Mechanism aimed toward Energy and Spectral Efficiency Gains	5
1.2.3 Extensions to Simulation Environments for Small Cell LTE Networks	7
1.3 Chapter outline	7
1.4 Notation	8
1.5 Publications	10
1.5.1 Peer Reviewed Publications	10
1.5.2 Non-Peer Reviewed Publications	11
2 Background and Related Work	13
2.1 Advanced Multi-Antenna Capabilities	14
2.1.1 Single-User MIMO (SU-MIMO)	15

2.1.2	Multi-User MIMO (MU-MIMO)	21
2.1.3	Coordinated Multi-Point (CoMP)	29
2.2	Advances in Small cells	35
2.2.1	The Use of Multiple Antennas in Small Cells	36
2.2.2	Collaboration and Supporting Network Structures	38
2.2.3	Green Wireless	40
2.3	Summary	41
3	Simulation Methodology and Scenarios	43
3.1	Comparison of Simulators	44
3.1.1	Simulator Types	45
3.1.2	Supported Features	48
3.1.3	Comparative Analysis	54
3.2	Extensions to the Vienna SL Simulator	59
3.2.1	Multi-User MIMO System Level Implementation	60
3.2.2	MU-MIMO User Selection Algorithm	62
3.2.3	Zero-Forcing MU-MIMO	64
3.2.4	MU-MIMO CQI calculation	65
3.2.5	Neighbouring eNB Feedback	66
3.2.6	MU-MIMO Spectral Efficiency	67
3.2.7	MMSE-IRC Filtering	68
3.2.8	Reassignment	68
3.3	Indoor and Outdoor Scenarios	69
3.3.1	Indoor Residential Scenario	70
3.3.2	Outdoor City Centre Scenario	71
3.4	Summary	73
4	MU-MIMO Across Small Cells: Spectral Efficiency Increases	77
4.1	Introduction	77
4.2	System Model	80
4.3	Reassignment Mechanism	83
4.3.A	Selection of Considered UEs	83
4.3.B	Check for Target UEs	85
4.3.C	Selection of UE to reassign from the set of Reassignable UEs	86

4.3.D	Reassignment	86
4.4	Simulation Results	88
4.4.1	Reassignability of UEs	91
4.4.2	Increases in MU-MIMO usage due to UE reassignments	94
4.4.3	Increases in Spectral Efficiency due to UE Reassignments	99
4.4.4	Time Scale Considerations	104
4.5	Summary	106
5	MU-MIMO Across Small Cells: Energy and Spectral Efficiency Increases	107
5.1	Introduction	107
5.2	System Model	111
5.3	Reassignment Mechanism	113
5.3.A	Selection of Considered UEs	113
5.3.B	Check for Target UEs	115
5.3.C	Selection of UEs to reassign from the Reassignable Set	115
5.3.D	Reassignment	116
5.3.E	Deactivation of emptied small cells	116
5.4	Simulation Results	117
5.4.1	Energy Efficiency Gains	118
5.4.2	Increases in Spectral Efficiency	123
5.4.3	Discussion of Other Potential Baselines	126
5.5	Summary	129
6	Conclusions and Future Work	131
6.1	Contributions and Findings	131
6.2	Future Work	134
6.2.1	Investigation of Additional Baselines for Ch5 Comparison	134
6.2.2	Dynamic Investigation of Ch5's Long-term Performance	135
6.2.3	Variation of Tx Power with Small Cell Density	135
6.2.4	Coordinated Reassignments between MEA-equipped Small Cells	135
6.2.5	Making Simulator Code Available	136
6.2.6	Combined MU-MIMO Across Small Cells and CS-CoMP	136
	Acronyms	137

Bibliography

143

List of Figures

1.1	MU-MIMO-based reassignment example.	4
1.2	Example outcome of reassignment resulting in small cell sleep state usage.	6
2.1	Downlink Multi-antenna transmission techniques.	14
2.2	Single-User MIMO transmission.	16
2.3	Comparison of SU-MIMO transmissions modes at UE speed of $3kmph$	18
2.4	Comparison of SU-MIMO transmissions modes at UE speed of $120kmph$	19
2.5	LTE Release 8 single-layer codebook beams for 4 Tx antenna uniform linear array.	20
2.6	Multi-User MIMO transmission.	22
2.7	Percentage throughput gain from using MU-MIMO over SU-MIMO.	23
2.8	Correlations between LTE Rel.8 codebook 4 Tx single-layer precoding vectors.	24
2.9	ZFBBF 2×1 computational complexity comparison of MU-MIMO schedulers.	27
2.10	ZFBBF 2×1 spectral efficiency comparison of MU-MIMO scheduling algorithms.	28
2.11	CoMP transmission Types.	30
2.12	LTE-A Downlink Reference Signal Structure for 2 Tx antenna configuration.	33
3.1	SINR distribution in an extremely dense heterogeneous deployment.	45
3.2	Mapping between simulator types and LTE protocol stack layers.	46
3.3	Curves used for mapping Effective SINR to BLER for a given CQI.	49
3.4	Variation in simulation time with small cell deployment density.	57
3.5	Dual Stripe indoor femtocell deployment scenario.	70
3.6	3-cell Munich outdoor picocell deployment scenario	72
3.7	21-cell Munich outdoor picocell deployment scenario.	74
4.1	Spectral Efficiency-based MU-MIMO Across Small Cells reassignment process	84
4.2	Spectral efficiency reassignment gains for varying Δ_{MUI}	88
4.3	Reassignment Explanatory Example.	91
4.4	Percentage of UEs that are reassignable.	93

4.5	Increases in Multi-User MIMO usage resulting from reassignment.	96
4.6	Multi-User MIMO Usage Before and After Reassignment.	97
4.7	Increases in Multi-User MIMO usage resulting from reassignment for different DRs	98
4.8	Increases in Spectral Efficiency resulting from reassignment.	100
4.9	Spectral Efficiency Before and After Reassignment.	101
4.10	Increases in Spectral Efficiency resulting from reassignment for different DRs. . .	102
4.11	Spectral Efficiency Before and After Reassignment for different DRs.	103
4.12	MU-MIMO usage resulting from reassignment for different time periods.	105
5.1	Small cell deactivation after MU-MIMO-based small cell reassignments.	110
5.2	MU-MIMO Across Small Cells with sleep mode usage mechanism.	112
5.3	Percentage of eNBs that can be deactivated.	120
5.4	Reduction in power consumption.	122
5.5	Deactivated eNBs for Greedy heuristic and Optimal reassignment selection. . . .	123
5.6	Increases in Spectral Efficiency resulting from reassignment.	124
5.7	Increases in Spectral Efficiency CDF.	125

List of Tables

3.1	Ease-of-use Metric/Indicators	54
3.2	Simulation Parameters and Hardware Configuration	58
3.3	SINRs used in MU-MIMO Simulation (single cell implementation) . .	61
4.1	Simulation parameters	92
4.2	Cell reassignment observations	94
5.1	Simulation parameters	119
5.2	Association Approaches for Small Cells with Multiple Antennas . . .	128

1 Introduction

This thesis shows that simple collaboration between base stations in small cell scenarios can lead to enhanced multiple-antenna gains. We demonstrate this for the case of *Multi-User MIMO-based reassignments* between adjacent small cells, from the perspectives of spectral efficiency, energy efficiency and number of spatial layers effectively utilised.

1.1 Overview

The use of multiple antennas in cellular systems (also known as Multiple Input Multiple Output (MIMO)) has recently led to significant gains in capacity and is commonly considered the key technological advance to which the transition from 3G to 4G was attributed. These gains come in three main forms: diversity, beamforming, and spatial multiplexing. Diversity gains improve the robustness of a link by transmitting and/or receiving different versions of the same signal over multiple antennas. Beamforming gains are obtained by applying complex weights to closely-spaced antennas in order to produce constructive and destructive interference in desired directions, hence boosting the Signal to Interference and Noise Ratio (SINR). Spatial multiplexing takes advantage of multipath fading effects between multiple Transmit (Tx) and Receive (Rx) antennas to simultaneously transmit multiple spatial streams which can be successfully separated at the receive side; in this way the throughput can be greatly elevated through the usage of an increased number of independent data streams. Of these, spatial multiplexing has attracted the most interest due to its ability to boost capacity so that it scales with the minimum of the number of Tx antennas n_t and

Rx antennas n_r .

This spatial multiplexing operation is not ideal, though, as correlations between the paths of different Tx-Rx pairs can lead to what are known as rank deficiencies, in which it is not possible to distinguish between transmissions on different spatial layers, leading to a reduction in throughput. Further, while it is often possible to scale up the number of transmit antennas at the base station side, due to the small form factor of the average mobile device, the housing of many antennas at the User Equipment (UE) side is more difficult. This can, in turn, limit the number of spatial layers available.

For these reasons there has been a recent shift in emphasis from spatial multiplexing on a single link between a base station and a UE, to what is known as Multi-User MIMO (MU-MIMO) [1], this was also forecasted in [2]. In MU-MIMO, instead of serving a single UE on multiple spatial layers, multiple UEs are each served on a single spatial layer, which is significantly less susceptible to rank deficiencies and allows the capacity to instead increase to the order of $\min(n_t, \sum_{k=1}^K n_{r,k})$, where $k \in \{1, \dots, K\}$ are the simultaneously served UEs within the cell. Essentially this consists of performing beamforming toward multiple spatially distinct UEs at once in such a way that the interference between the beams directed at co-scheduled UEs is kept as low as possible.

In this thesis we focus on the use of MU-MIMO in coordinated Long Term Evolution (LTE) small cell networks. Small cell infrastructure has recently seen massive growth in deployment with the number worldwide currently exceeding 8.4 million [3] (compared to roughly 6 million macrocells [4]). This is a very appealing technology to operators due to its low cost of deployment and operation compared to macrocell base stations, and its ability to dramatically increase the spatial reusability of limited spectrum.

MU-MIMO works best when the antennas at the base station experience highly correlated channel characteristics [2, Fig.4]; as high antenna correlation is often observed in cases where antennas are closely spaced this suits the small form factor of small cell devices. Further, due to strong line-of-sight components and reduced scattering in small cell scenarios compared to macrocell cases [5], spatial multiplexing on a single base station-to-UE link would be subject to regular rank deficiencies. MU-MIMO does however, for effective operation, require that the channels between the base station and different simultaneously-served UEs be

sufficiently uncorrelated (close to orthogonal) to ensure that excessive cross-layer interference is not experienced. While this cross-layer or multi-user interference can be suppressed at the transmitter using techniques such as Zero-Forcing Beam Forming (ZFBBF), there is an associated cost in terms of signal power, meaning that the combined throughput of the spatial layers of multiple non-orthogonal UEs may be less than the unsuppressed throughput of a single UE transmitted on a single layer alone. In macrocell networks this problem is overcome by taking advantage of multi-user diversity through MU-MIMO scheduling. In other words, as macrocells contain many UEs, user selection can ensure that sets of UEs with mutually orthogonal channels are always scheduled together. On the other hand, in small cell scenarios where the number of UEs per cell is low, sets of UEs with mutually orthogonal channels do not always exist in a given cell and the additional spatial multiplexing gains obtained for MU-MIMO cannot always be realised. At the same time, due to the high density of deployment in small cell networks, UEs are often in range of multiple small cells at once.

For this reason we investigate the reassignment of UEs between adjacent small cells to increase both the channel orthogonality between the UEs of each cell in the small cell network and the number of spatial layers which can be effectively utilised. We propose a number of network-wide mechanisms which target different performance objectives through the use of centralised coordination between neighbouring small cells. These objectives include increases in spectral efficiency and reductions in small cell network energy consumption through combined use of MU-MIMO-based reassignments and small cell sleep states. We collectively call these mechanisms *Multi-User MIMO across Small Cells*.

In investigating our proposed *Multi-User MIMO across Small Cells* mechanisms we compare to a primary baseline of Reference Signal Received Power (RSRP)-based cell association. This represents the default association policy in LTE systems [6, Section 5.2.3] and, in cases where all evolved Node Base stations (eNBs) operate within the same frequency and have the same multi-antenna capabilities, this is the association scheme which maximises the UE's average SINR [7].




Figure 1.1: MU-MIMO-based reassignment example. UE colours indicate the small cell base station to which each UE is initially attached.

1.2 Key Contributions

In devising new methods to enhance downlink multi-antenna transmission capabilities through small cell collaboration the following primary contributions are made:

1.2.1 MU-MIMO-based Reassignment Mechanism aimed toward Spectral Efficiency Gains

This contribution centres on the proposal and assessment of the *Multi-User MIMO across Small Cells* concept, from the perspective of spectral efficiency gains. As already stated, MU-MIMO is the spatial multiplexing of data streams to multiple UEs, each on different spatial layers, which can considerably increase UE spectral efficiencies. In small cell networks, where the number of UEs per cell is low, finding suitable sets of UEs to be co-scheduled for MU-MIMO is not always possible. In these cases we propose MU-MIMO-based cell reassignments of UEs into adjacent cells to engineer a scenario in which the number of sets of UEs with relative channels suitable for MU-MIMO operation is increased. This work investigates the gains of

this new concept for realistic indoor and outdoor small cell scenarios, taking into account user scheduling before and after the reassignment occurs. From system level simulations we show that the proposed mechanism results in considerable increases in spectral efficiency and MU-MIMO usage by both the reassigned UE and the UE in the target cell with which it gets paired. We also show that, in the outdoor small cell scenario, the percentage of UEs that meet the reassignment criteria and can benefit from neighbour-cell MU-MIMO is almost double that of the indoor scenario.

In this contribution the emphasis is purely on increased spectral efficiency. Let us consider the simple scenario shown in Figure 1.1 where UE colours indicate the small cell eNBs to which the UEs are initially attached and the marked Precoding Matrix Indicator (PMI) of each UE indicates a fed-back quantised form of the channel directionality of each eNB-to-UE link. To simplify explanation, this figure assumes that the PMI of each UE stays the same regardless of which eNB it is served by; in reality, the channels from different eNBs are likely to instead correspond to different PMIs. As will be illustrated later in Figure 2.8 the channel codeword corresponding to PMI=9 is orthogonal to those of both PMI=1 and PMI=10, while other PMI codeword combinations are not orthogonal. Assuming that only UEs with orthogonal quantised channels can be co-scheduled for MU-MIMO, MU-MIMO cannot be performed in any small cell in this configuration. However, by reassigning either UE1 to eNB 1 or UE 2 to eNB 2, assuming that UE SINRs after reassignment are sufficiently high in both cases, previously unrealisable MU-MIMO spatial multiplexing gains can be obtained.

1.2.2 Refined MU-MIMO-based Reassignment Mechanism aimed toward Energy and Spectral Efficiency Gains

In this contribution we investigate the reassignment of UEs between adjacent small cells to concurrently enable spatial multiplexing gains through MU-MIMO and reductions in energy consumption through switching emptied small cells to a sleep state. As before, UEs can be reassigned between adjacent small cells provided that the resulting UE-to-base station assignment corresponds to increased mutual orthogonality between the UEs served within each cell and a minimum expected increase in the spectral efficiency of the reassigned UE over the earlier




Figure 1.2: Example outcome of reassignment resulting in small cell sleep state usage.

configuration. We formulate the selection decision of which UEs to reassign as a set covering problem with the objective of maximising the number of small cell base stations to switch to a sleep state. Our results show that, for both indoor and outdoor LTE small cell scenarios, the proposed MU-MIMO-based reassignments can achieve significant reductions in the required number of active small cell base stations, whilst simultaneously achieving increases in spectral efficiency of the affected UE set for which MU-MIMO is enabled (i.e. the set of the reassigned UE and the target UE with which they get paired in the target cell).

Returning to the example in Figure 1.1, two possible reassignments exist, both of which result in a configuration better suited to MU-MIMO spatial multiplexing gains. These are the reassignment of UE1 to eNB1 and the reassignment of UE2 to eNB2. It is also the case that performing both reassignments does not make sense as if UE2 is reassigned to eNB2 then it cannot perform MU-MIMO with UE1 in eNB1, as each UE can only be attached to a single eNB. In the previous mechanism the decision of which of these two reassignments to perform was based purely on which of the two reassignments results in the greatest expected spectral efficiency increase for the reassigned UE. In the refined reassignment mechanism, however, priority is given

to the reassignment of UE2 to eNB2 as this results in the emptying of eNB1. If there are no longer any UEs attached to eNB1 it can be switched to a sleep state and in doing so both conserve energy and reduce interference to the UEs in the neighbouring cells. This is illustrated in Figure 1.2.

1.2.3 Extensions to Simulation Environments for Small Cell LTE Networks

In the course of implementing the two above investigations for system-level simulation a number of contributions from the point of view of simulation methodology were made. In order to assist others in the simulation of Extremely Dense Networks (EDNs) we identify and compare a representative cross-sectional set of potential EDN simulators and put forward guidelines to steer readers toward the selection of the appropriate simulator for their desired investigation. Further, in the context of the Vienna System Level (SL) simulator [8], which we have selected as the baseline simulator for our investigations, a number of extensions are required in order to assess our desired scenarios. These extensions include, but are not limited to, the implementation of MU-MIMO at system level, MU-MIMO scheduling and channel quality estimation, and the inclusion of indoor and outdoor small cell scenarios. We provide discussion of the extensions performed and how they are carried out, and we will make some of these extensions available to the wider community so that they may be built upon by others.

1.3 Chapter outline

The remainder of the thesis is organised as follows:

As our work considers the reassignment of UEs between adjacent small cells to achieve increased multi-antenna gains, *Chapter 2* focuses on providing background on the state of the art in the fields of advanced multi-antenna techniques and small cell deployments, including a discussion of multi-cell coordination within both contexts, how the two fields cross over, and, in the case of small cells, how they can be advanced toward more energy efficient operation.

To assess the potential gains of our proposed reassignment mechanisms it is essential to be able to simulate their performance in the context of realistic scenarios. How we go about this is the subject matter of *Chapter 3*. We provide a comparative analysis of simulation tools suitable for densely deployed small cell network investigations and compile guidelines for others wanting to select an appropriate simulation environment. Based on this comparison, for our investigations we take as a baseline the Vienna LTE SL Simulator. We then identify, implement, and detail the required extensions to this simulation environment for our investigated case.

Chapter 4 focuses on our contributions obtained from the *MU-MIMO Across Small Cells* reassignment mechanism, primarily from the perspectives of spectral efficiency and MU-MIMO spatial multiplexing gains. The performance of our MU-MIMO-based reassignment mechanism is assessed in both realistic indoor and outdoor environments and with varying deployment densities.

Chapter 5 extends the work of Chapter 4 by also considering the application of small cell sleep states to reduce the amount of energy consumed by the small cell network. In order to maximise the number of small cells switched to a sleep state while ensuring that reassignments result in an increase in the reassigned UE's spectral efficiency, we reformulate the reassignment selection process into the form of a set covering problem. We then present the increases in both energy and (reassigned and target UE) spectral efficiency that are achieved from our proposed approach.

Chapter 6 concludes the work by summarising our findings and identifying a number of directions for future work.

1.4 Notation

Throughout this thesis we apply the following conventions:

- Italic letters (e.g., h) represent *scalars*.
- Lower case boldface letters (e.g., \mathbf{h}) represent *vectors*.
- Upper case boldface letters (e.g., \mathbf{H}) represent *matrices*.
- Calligraphic letters (e.g., \mathcal{U}) represent *sets*.

Additionally, the following listed variables are used regularly:

- \mathbf{H} , \mathbf{h} or h denote (matrices or vectors of) channel coefficients.
- \mathbf{x} or x denote transmitted signals, before transmit precoding.
- \mathbf{y} or y denote received signals, before receive filtering.
- \mathbf{z} or z denote received signals, after receive filtering.
- \mathbf{U} or \mathbf{u} denote codebook-defined fed back unitary precoders.
- \mathbf{W} or \mathbf{w} denote transmitted unitary precoders.
- \mathbf{G} or \mathbf{g} denote receive filters.
- \mathbf{h}_{eff} denotes effective channels (product of the channel matrix \mathbf{H} and precoding vector \mathbf{w}).
- \mathbf{C} denotes covariances.
- \mathbf{n} denotes noise.
- k and j denote user indices (general and interfering, respectively).
- e , l , O and T denote base station indices (general, interfering, serving/original and target base stations, respectively).
- b denotes the index of a Resource Block (RB)(subband).
- n_{MU} denotes the number of users simultaneously served by a MU-MIMO transmission, with a maximum value of N_{MU} .
- r denotes the predicted instantaneous rate.
- R denotes the long-term average rate.
- a indicates whether a small cell base station is operating in an active state.

$\|\cdot\|$, $(\cdot)^{-1}$, $(\cdot)^H$, $(\cdot)^T$ and $(\cdot)^\dagger$ denote the norm, inverse, conjugate (Hermitian) transpose, regular transpose and Moore-Penrose pseudo-inverse of a matrix, respectively. $|\cdot|$ denotes the absolute value of scalar. \emptyset represents the null/empty set. Finally, all uses of the term *average* refer to the mean value.

1.5 Publications

Publications which relate directly to the work of this thesis are marked with a black bullet symbol (●). Publications marked with a white bullet symbol (○) were performed as part of other projects during the course of this PhD.

1.5.1 Peer Reviewed Publications

Submitted

- D. Finn, H. Ahmadi, R. Razavi, H. Claussen, and L. A. DaSilva, "Energy and Spectral Efficiency Gains From Multi-User MIMO-based Small Cell Reassignments," submitted to *IEEE Global Communications Conference (Globecom)*, 2015.
- D. Finn, H. Ahmadi, A. Cattoni, and L. A. DaSilva, "Improved Spectral Efficiency through Multi-User MIMO Across Small Cells," submitted to *IEEE Transactions on Vehicular Technology (TVT)*, 2015.
- D. Finn, C. Galiotto, P. Alvarez, J. Van De Belt, H. Ahmadi, and L. A. DaSilva, "Simulating Dense Small Cell Networks," submitted to *IEEE Vehicular Technology Magazine*, 2014.

Accepted

- D. Finn, H. Ahmadi, A. Cattoni, and L. A. DaSilva, "Multi-User MIMO across Small Cells," in *IEEE International Conference on Communications (ICC)*, Sydney, Australia, 2014.
- I. Macaluso, D. Finn, B. Oztgul, and L. A. DaSilva, "Complexity of Spectrum Activity and Benefits of Reinforcement Learning for Dynamic Channel Selection," in *IEEE Journal on Selected Areas in Communications (JSAC)*, vol. 31, no. 11, pp. 2237-2248, Nov. 2013.

- P. Van Wesemael, W. Liu, M. Chwalisz, J. Tallon, D. Finn, Z. Padrah, S. Pollin, S. Bouckaert, I. Moerman, and D. Willkomm, "Robust Distributed Sensing with Heterogeneous Devices," in *Future Network & Mobile Summit (FNMS)*, Berlin, Germany, 2012.
- D. Finn, J. Tallon, L. A. DaSilva, S. Pollin, W. Liu, S. Bouckaert, and J. V. Gerwen, "Experimental Assessment of Tradeoffs among Spectrum Sensing Platforms," *Sixth ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation and Characterization (WiNTECH)*, Las Vegas, Nevada, USA, 2011.

1.5.2 Non-Peer Reviewed Publications

- D. Finn, H. Ahmadi, A. F. Cattoni, and L. A. Dasilva, "Multi-user MIMO in small cell networks with coordinated scheduling," in *4th Workshop of COST Action IC0902*, Rome, Italy, 2013.
- D. Finn, L. DaSilva, and Y. Xiao, "Playing Games in Cognitive Radio Networks," *IEEE Communications Society Multimedia Communications Technical Committee E-Letter (COMSOC MMTC E-Letter)*, vol. 7, no. 6, pp. 22-25, 2012.
- J.-H. H. C. Heller, S. Bouckaert, I. Moermann, S. Pollin, P. v. Wesemael, D. Finn, D. Willkomm, "A Performance Comparison of Different Spectrum Sensing Techniques.", *Wireless Innovation Forum Conference on Wireless Communications Technologies and Software Defined Radio (SDR-WInnComm)*, 2011.
- L. A. Dasilva, L. Doyle, D. Finn, and J. Tallon, "CREW: Building a Cognitive Radio Federation," in *1st Workshop of COST Action IC0902*, Bologna, Italy, 2010.

2 Background and Related Work

In this chapter we provide background and related work on two main concepts which are central to this thesis: advanced multi-antenna techniques and small cell deployments.

We start this by providing some background on single-cell downlink multi-antenna concepts: *what they are, how they work* and *what general directions the research is going in*, before moving on to coordinated multi-cell concepts and discussing how our work ties in with these.

Next we consider small cell deployments, where we define small cell following the definition of the Small Cell Forum [9]: "'Small cells' is an umbrella term for operator-controlled, low-powered radio access nodes, including those that operate in licensed spectrum and unlicensed carrier-grade Wi-Fi. Small cells typically have a range from 10 meters to several hundred meters. Types of small cells include femtocells, picocells and microcells - broadly increasing in size from femtocells (the smallest) to microcells (the largest). Any or all of these small cells can be based on 'femtocell technology' - i.e. the collection of standards, software, open interfaces, chips and know-how that have powered the growth of femtocells."; however, we include Remote Radio Heads (RRHs) as a type of small cell also. In this chapter, we consider how the small cell market is progressing, what multi-antenna techniques have been investigated for these scenarios, what architectures are best for allowing coordination between small cells, and finally, how further energy savings can be made in small cell green wireless networks.




Figure 2.1: Downlink Multi-antenna transmission techniques.

2.1 Advanced Multi-Antenna Capabilities

Multiple Input Multiple Output (MIMO) uses multiple transmit and receive antennas to improve channel performance. In this thesis we centre our focus on downlink multi-antenna transmissions, which, as shown in Figure 2.1, can be broadly broken down into three main categories: Single-User MIMO (SU-MIMO), Multi-User MIMO (MU-MIMO) and Coordinated Multi-Point (CoMP) transmissions. These correspond, respectively, to transmissions from a single multi-antenna transmitter to a single multi-antenna receiver, transmissions from a single multi-antenna transmitter to a multiple (either single- or multi- antenna) receivers (each on separate spatial layers), and coordinated transmissions from multiple transmitters to a single (or multiple) receiver(s). Within these categories a number of sub-categorisations also exist.

In this thesis we will focus mainly on 3rd Generation Partnership Program (3GPP) Long Term Evolution (LTE)/LTE-Advanced (LTE-A) systems; however, the mechanism we propose could equally be applied to other radio access technologies.

Right from its initial release in 2008, LTE standardisation has relied on the use of multiple antenna transmissions to boost system performance, supporting up to four antennas at both the Transmit (Tx) and Receive (Rx) side and transmissions on up to four spatial-layers. This initial release focused on SU-MIMO transmissions and offered only basic support for MU-MIMO. Since then, increased emphasis was placed on MU-MIMO capabilities with the introductions in LTE-A Rel.10 of

transmission mode 9 and the Demodulation Reference Signal (DM-RS). As well as enabling SU-MIMO transmissions on up to 8 spatial layers, these allowed for flexible switching between SU-MIMO and MU-MIMO, non-codebook-based precoding and MU-MIMO spatial multiplexing of up to 4 User Equipments (UEs) each on a single spatial layer or 2 UEs each on 2 spatial layers. Since then MU-MIMO advances focusing on improving Channel State Information (CSI), 3D beamforming (MU-MIMO serving multiple UEs on the elevation plane) and massive MIMO (MU-MIMO with very large numbers of co-scheduled UEs) have been discussed for the Release 12 [1].

CoMP functionalities were first discussed for LTE Rel.9 although it wasn't until Rel.11 that basic CoMP support was provided. Continuing specification efforts in this area are primarily focused on improving multi-point CSI feedback specification, control channel enhancements and CoMP solutions which make use of relaxed backhaul requirements [1].

Needless to say, the current emphasis toward multi-antenna techniques which exploit network/system configurations (whether multiple UEs or multiple evolved Node Base stations (eNBs)) ties in well with our own emphases within this thesis.

2.1.1 Single-User MIMO (SU-MIMO)

In an SU-MIMO general system with N_t Tx antennas and $N_{r,k}$ Rx antennas, performed over N_L spatial layers, the $N_{r,k} \times 1$ received signal of UE k , \mathbf{y}_k , can be represented as follows

$$\mathbf{y}_k = \underbrace{\mathbf{H}_{k,O}\mathbf{d}_k}_{\text{Desired Signal}} + \underbrace{\sum_{l=1}^{N_{eNB}-1} a_l \mathbf{H}_{k,l}\mathbf{d}_l}_{\text{Inter-Cell interference}} + \underbrace{\mathbf{n}_{W,k}}_{\text{Noise}} \quad (2.1)$$

where $\mathbf{H}_{k,e}$ is the $N_{r,k} \times N_t$ channel matrix from eNB e to UE k , where e can be either O for the original/serving cell or l for interfering cells, and \mathbf{d}_k and \mathbf{d}_l represent the $N_t \times N_L$ desired and interfering transmitted signals, respectively. $\mathbf{n}_{W,k}$ represent complex Additive White Gaussian Noise (AWGN), the elements of which have zero mean and variance σ^2 . a_l is a boolean indicator, showing whether interfering cell l is actively transmitting. This is illustrated in Figure 2.2 for a downlink 2×2 system




Figure 2.2: Single-User MIMO transmission.

not subject to inter-cell interference, where $h_{tx,rx}$ are the channel coefficients from Tx antenna tx to Rx antenna rx which make up the channel matrix $\mathbf{H}_{k,O}$.

Commonly in SU-MIMO transmissions, three main methods are discussed: transmit diversity, spatial multiplexing and beamforming, which are distinguished by their different performance objectives. **Transmit diversity** operates by transmitting the same signal over multiple antennas each coded differently using either space-time block codes (also known as Alamouti codes [10]) or their adapted form space-frequency block codes. This has the effect of increasing the robustness of the channel to error. **Spatial multiplexing** is the transmission of multiple data streams over different spatial layers in such a way that they can be separated at the receiver. This has the effect of increasing the capacity of the channel by the order of the number of spatial layers used, which has a maximum of $\min(N_t, N_{r,k})$. Finally, **beamforming** applies complex weights to transmissions on different antennas in such a way as to combine constructively in the desired direction and destructively in directions which would cause interference to other users. This has the effect of increasing the desired signal gain, while at the same time decreasing the level of interference created. For each of these open- and closed- loop flavours exist, corresponding to those in which channel knowledge at the transmitter (other than for Modulation and Coding Scheme (MCS) selection) is not, and is, required, respectively. Detailed explanations of how these methods work can be found in [11,

Ch 5].

Of the SU-MIMO methods there is no one-size-fits-all best solution for all wireless scenarios, hence in the standards a range of different multi-antenna *transmission modes* are often defined. Famously, in trying to identify the optimal way to use of multiple antennas, Zheng and Tse [12] quantified the tradeoff region between link robustness provided by diversity and additional throughput provided by spatial multiplexing. However, in general the choice of transmission mode depends on a number of factors including: Signal to Interference and Noise Ratio (SINR), UE speed, channel correlation and traffic type.

Broadly speaking, in SU-MIMO, open-loop techniques should be used in high velocity scenarios in which channel feedback is less reliable; diversity techniques should be used in scenarios where high channel reliability is required, such as control channels; spatial multiplexing should be used in cases where there is high SINR and low correlation between the channels of different antennas; and beamforming should be used in cases where the channels are highly correlated, where increased SINR is required (e.g. on the cell edge) or where control is needed over the emitted interference.

To demonstrate these points, Figures 2.3 and 2.4 compare the throughputs of four LTE Transmission Modes (TMs) in low and high speed scenarios, respectively, with varying Signal to Noise Ratio (SNR). These TMs are single antenna transmissions (*SISO*), Transmit Diversity (*TxD*), Open-Loop Spatial Multiplexing (*OLSM*), and Closed-Loop Spatial Multiplexing (*CLSM*). These graphs were generated using the Vienna Link Level (LL) downlink simulator [8] for a single eNB and single UE with 2 Hybrid Automatic Repeat reQuest (H-ARQ) retransmissions, a channel feedback delay of $5ms$, and using the Typical Urban (TU) channel model [13] with low correlation between transmit antennas. As can be seen, the multi-antenna techniques show superior performance to SISO transmissions in all cases except for the TxD high SNR case, for which the losses in throughput from feedback overhead exceed the gains from channel reliability. We see that at low speed (Fig. 2.3) the Closed-Loop transmissions provide higher throughput than the other TMs for all SNRs, while at high speeds (Fig. 2.4) the Open-Loop transmissions perform best with TxD showing best performance at low SNR and OLSM at high SNR. The humps in this graph result from the use of H-ARQ, which, in cases of retransmissions, effectively reduces the channel coding rate while




Figure 2.3: Comparison of SU-MIMO transmissions modes at UE speed of 3kmph.

keeping the same modulation scheme. Hence, for increasing SNR, the humps correspond to 4QAM, 16QAM and 64QAM, respectively.

In this thesis we focus on low mobility scenarios in which the correlation between antennas is high (corresponding to close antenna spacings, low levels of scattering, and/or strong Line-Of-Sight (LOS) components). For this reason, going forward, use of the term SU-MIMO will mainly refer to beamforming operation, also known as *Closed-Loop Spatial Multiplexing (CLSM)* on a single spatial layer in LTE terminology, owing to the fact that in LTE this transmission mode operates in the same way as CLSM, with the exception that only a single spatial layer is used.

In LTE CLSM on a single spatial layer the signal model can be expressed as follows

$$\mathbf{y}_k = \underbrace{\mathbf{H}_{k,O} \mathbf{w}_k x_k}_{\text{Desired Signal}} + \underbrace{\sum_{l=1}^{N_{eNB}-1} a_l \mathbf{H}_{k,l} \mathbf{d}_l}_{\text{Inter-Cell interference}} + \underbrace{\mathbf{n}_{W,k}}_{\text{Noise}} \quad (2.2)$$




Figure 2.4: Comparison of SU-MIMO transmissions modes at UE speed of 120kmph.

where the transmitted signal (\mathbf{d}_k in Eqn. (2.1)) consists of the transmitted symbol x_k with unitary precoding \mathbf{w}_k applied. The unitary precoding is selected from a predefined codebook specific to the number of Tx antennas used. Figure 2.5 shows the corresponding beam patterns for each of the entries in the LTE Rel.8 precoding codebook [14, Table 6.3.4.2.3-2] when applied to a Uniform Linear Array (ULA) of dipole antennas with a half-wavelength inter-antenna spacing. As can be seen these beam patterns generally have multiple strong lobes and multiple zeros in order to allow interference restriction in multiple directions at once. The precoders in each subfigure are all mutually orthogonal.

SU-MIMO Channel Feedback

The LTE standards specify three types of Channel State Information (CSI) feedback for use in the scheduling of UEs and adaptive modulation and coding. These are the Channel Quality Indicator (CQI), Precoding Matrix Indicator (PMI) and Rank

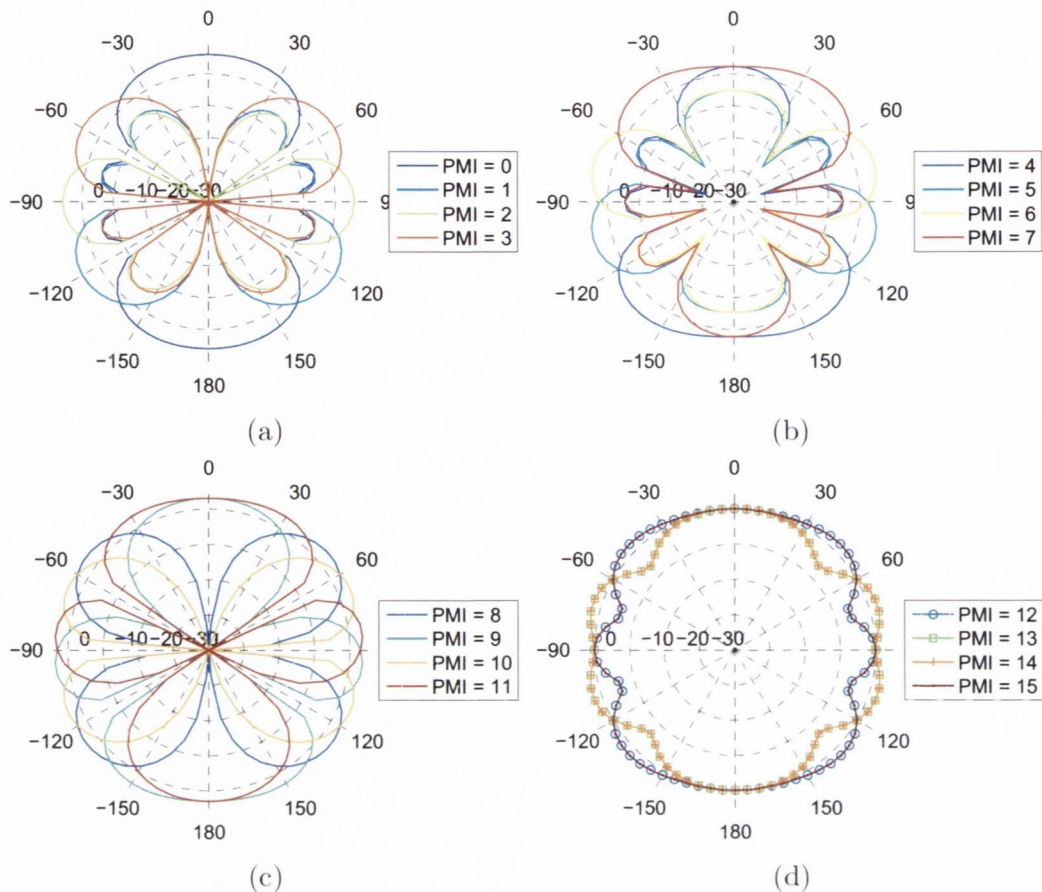


Figure 2.5: LTE Release 8 single-layer codebook beams for 4 Tx antenna uniform linear array.

Indicator (RI) [15]. The CQI is a form of SINR, quantised into one of 16 values, the PMI recommends a linear precoding matrix from a predefined codebook, and the RI indicates the rank (number of parallel transmission streams) to use. With regard to the RI, going forward in this work, for both SU-MIMO and MU-MIMO operation, only Rank-1 transmissions for each UE will be considered.

The CQI quantisation process proceeds as follows: for each of the 16 possible CQI values there is a corresponding MCS; the CQI is selected as the highest order MCS which is expected to achieve a Block Error Rate (BLER) no greater than 10% at the UE's SINR [16].

In LTE Rel.8 systems, recommended precoders are selected from a predefined codebook $\mathcal{C} = \{\mathbf{U}_0, \dots, \mathbf{U}_{2N_t-1}\}$ of $N_t \times N_L$ unit-norm matrices, and the index of the selected codeword in the codebook is called the PMI. Essentially this corresponds to a quantisation of the channel matrix, containing information

regarding directionality of the strongest channel components (or the strongest channel component for single-layer). The codebooks are defined in [14, Section 6.3.4]. PMI selection is based on the UE channel matrix although the metric used in the selection varies between implementations. For example, the metric may target the maximisation of the effective channel gain $I_b(\mathbf{U}_i) = \det(\mathbf{U}_{i_k}^H \mathbf{H}_{k,O,b}^H \mathbf{H}_{k,O,b} \mathbf{U}_{i_k})$ (or equivalently $I_b(\mathbf{U}_i) = |\mathbf{H}_{k,O,b} \mathbf{u}_{i_k}|^2$ for single-layer [17]), where $I_b(\mathbf{U}_i)$ is the metric computed for Resource Block (RB) (/subband) $b \in \{1, \dots, B\}$ and B is the number of RBs in the system bandwidth.

Both CQI and PMI feedback can be specified as either *wideband* or *subband*, while for the RI only wideband feedback is allowed. Put simply, wideband CSI corresponds to the best CSI, on average, if applied to the entire considered bandwidth, while subband CSI is separately specified for each subband RB.

Wideband CQI is computed by the UE by averaging the SINRs over all RBs using Mutual Information Effective SINR Mapping (MIESM) (or a similar Effective SINR Mapping (ESM) technique, as will be discussed in Section 3.1.2) before quantising the resulting effective SINR into one of the 16 possible values.

The wideband PMI is selected as the PMI, which, when applied to the entire considered bandwidth, achieves the highest sum metric. In other words [18]

$$\mathbf{U}_{\text{PMI}} = \underset{\mathbf{U}_i \in \mathcal{C}}{\operatorname{argmax}} \sum_{b=1}^B I_b(\mathbf{U}_i). \quad (2.3)$$

In LTE-A Rel.10 systems for 4 Tx antennas or less this is performed exactly the same as LTE Rel.8, using the exact same codebook. For more than 4 Tx antennas a new dual codebook is defined in which *both* wideband and subband PMIs are selected and fed back. In our investigations later in this work 4 Tx antennas will be used, meaning that the LTE Rel.8 codebook is utilised.

The feedback specification methods discussed in this subsection apply to both SU-MIMO and MU-MIMO transmissions.

2.1.2 Multi-User MIMO (MU-MIMO)

MU-MIMO is a spatial multiplexing technique in which multiple transmit antennas at the eNB are used to simultaneously serve multiple UEs within a single time frequency




Figure 2.6: Multi-User MIMO transmission.

resource as shown in figure 2.6; this is achieved by transmitting to each UE on a different spatial layer. Essentially this consists of applying precoding weights to direct orthogonal beams at each served UE; in other words, performing beamforming with multiple UEs at once in such a way that the interference between co-scheduled UEs is minimised. Due to MU-MIMO's ability to provide large spatial multiplexing gains without requiring additional antennas on the UE, and its ability to overcome rank deficiency problems (which often limit single point-to-point spatial multiplexing), MU-MIMO capabilities have been highly emphasised in recent 3GPP standardisation and in research projects such as FP7 Spectrum Aggregation and MU-MIMO: Real-world Impact (SAMURAI) [19].

In a downlink scenario where a base station with N_t transmit antennas transmits to K UEs, each with $N_{r,k}$ receive antennas, the capacity of the channel scales with $\min(N_t, \sum_{k=1}^K N_{r,k})$. There are, however, some additional losses in the SINR, firstly, due to the splitting of the transmit power across the co-scheduled UEs' transmissions, secondly due to any orthogonalisation operations (e.g. Zero-Forcing Beam Forming (ZFBE)) in the precoding of transmissions to UEs with non-orthogonal quantised channels, and thirdly, due to any residual Multi-User Interference (MUI) between the UEs.

In order to minimise MUI and ZFBE losses, only UEs which have fed back close




Figure 2.7: Percentage throughput gain from using MU-MIMO over SU-MIMO.

to orthogonal (semi-orthogonal) quantised channel vectors (as specified by the PMI) should be co-scheduled for MU-MIMO transmissions. Any residual MUI, which primarily results from channel quantisation errors, should then be suppressed by the receiver.

As an illustration of what MU-MIMO gains are possible, in Figure 2.7 we plot the percentage throughput increases of MU-MIMO, with two co-scheduled UEs, over Single-User MIMO beamforming on a single spatial stream (SU-MIMO), for different levels of residual MUI after suppression, represented by the factor Δ_{MUI} . In this figure both UEs served simultaneously using MU-MIMO have equal SINRs and interference suppression capabilities, and their quantised channels are orthogonal. Δ_{MUI} ranges from 0 to 1, with higher Δ_{MUI} representing higher residual MUI (poor MUI suppression) and $\Delta_{MUI} = 0$ representing no residual MUI (full MUI suppression). It must also be stated that due to CQI feedback quantisation errors the achieved gains will be slightly less than those shown, while




Figure 2.8: Correlations between LTE Rel.8 codebook 4 Tx single-layer precoding vectors.

gains at very low SINR are lost. As a result, a UE can benefit from MU-MIMO transmissions if, firstly, its SINR is sufficiently high (greater than roughly 1.5dB), and secondly, there is another UE with a precoding matrix semi-orthogonal to its, with which to pair.

Figure 2.7 was generated by computing 2-layer MU-MIMO SINRs for the given range of single-layer SU-MIMO SINRs using Eqn.(3.6), which is obtained from [20]. Shannon rates are computed for these SINRs and, in the case of MU-MIMO, summed over the two co-scheduled UEs.

As mentioned, in LTE and LTE-A systems, the orthogonality of co-scheduled UEs is assessed based on their fed back PMI which indicates the directionality of the channel between the base station and the UE in quantised form. To illustrate the orthogonality between different LTE Rel.8 codebook precoders, Figure 2.8 displays the correlation of each possible 2 UE precoder combination in the 4 Tx antenna case. If the correlation is zero (shown in dark blue) the codewords are orthogonal, increased correlation corresponds to decreased orthogonality. We observe that, in this

case, each codeword is orthogonal to at least 3 other codewords and that mutually orthogonal codewords tend to be indexed in blocks of 4.

MU-MIMO Signal Model

For MU-MIMO with n_{MU} UEs co-scheduled, each transmitted to on a single spatial layer (rank-1), the received signal of UE k , co-scheduled for MU-MIMO with UEs $j \in \{1, \dots, n_{MU} - 1\}$, is given by

$$\mathbf{y}_k = \underbrace{\mathbf{H}_{k,O} \mathbf{w}_k x_k}_{\text{Desired Signal}} + \underbrace{\sum_{j=1}^{n_{MU}-1} \mathbf{H}_{k,O} \mathbf{w}_j x_j}_{\text{Multi-User Interference}} + \underbrace{\sum_{l=1}^{N_{eNB}-1} a_l \mathbf{H}_{k,l} \mathbf{W}_l \mathbf{x}_l}_{\text{Inter-Cell interference}} + \mathbf{n}_{W,k} \quad (2.4)$$

where \mathbf{y}_k represents the $N_r \times 1$ received signal vector, $\mathbf{H}_{k,O}$ represents the $N_r \times N_t$ channel matrix within the original serving cell, \mathbf{w}_k represents the $N_t \times 1$ applied unitary precoding and x_k represents the transmitted symbol, of UE k . $\mathbf{H}_{k,O} \mathbf{w}_j x_j$ represents the interference from co-scheduled UE j , while $\mathbf{H}_{k,l} \mathbf{W}_l \mathbf{x}_l$ represents the interference from neighbouring cell l and finally $\mathbf{n}_{W,k}$ represent complex AWGN the elements of which have zero mean and variance σ^2 . It should be noted that the value $\mathbf{H}_{k,O}$ in this equation includes the transmit power to each of the co-scheduled UEs, which is $\frac{1}{n_{MU}}$ of the power it would have if they were scheduled for SU-MIMO alone. As indicated by a_l , if a neighbouring eNB is not active it does not produce inter-cell interference. It is assumed that all serving cells are active, hence we omit the a_0 which would otherwise precede $\mathbf{H}_{k,O}$.

After a $1 \times N_r$ receive filter \mathbf{g}_k is applied we get the received symbol z_k of user k as

$$z_k = \mathbf{g}_k \mathbf{y}_k. \quad (2.5)$$

The post-reception SINR of UE k on a given subcarrier can be represented as

$$\gamma_k = \frac{|\mathbf{g}_k \mathbf{H}_{k,O} \mathbf{w}_k|^2}{|\mathbf{g}_k \sum_{j=1}^{n_{MU}-1} \mathbf{H}_{k,O} \mathbf{w}_j|^2 + |\mathbf{g}_k \sum_{l=1}^{N_{eNB}-1} a_l \mathbf{H}_{k,l} \mathbf{W}_l|^2 + \sigma^2 \mathbf{I} \|\mathbf{g}_k\|^2}. \quad (2.6)$$

MU-MIMO Transmission Strategies

MU-MIMO transmissions can be performed in a number of different ways. *Dirty Paper Coding (DPC)* [21] is a non-linear scheme which provides an upper bound on MU-MIMO performance but is highly complex and requires full knowledge of the interference signal when decoding, making it not practical in reality. Instead this is often used as a capacity upper bound for MU-MIMO channels. *Zero-Forcing Beam Forming (ZFBF)* [22] is a linear precoding technique in which UEs are each served on a single spatial layer and orthogonal precoders are selected for each UE in an attempt to nullify MUI. *Block Diagonalisation (BD)* [23] is another linear precoding technique, this time with UEs each served on multiple spatial layers. This method operates by nullifying interference between different UEs, although not nullifying the interference between layers sent to the same UE. In this way a reduction in the power losses required to enforce orthogonality can be achieved as interference between layers sent to a single UE can be successfully removed at the UE receiver. *Zero-Forcing (ZF) DPC* [22] combines Zero-Forcing Beam Forming (ZFBF) and Dirty Paper Coding (DPC) to achieve near optimal performance with much reduced complexity (compared to DPC). This method however requires non-causal channel knowledge and so cannot be used in real systems. *Successive ZF DPC* [24] is similar to BD although for the ZF DPC case.

In LTE Rel.8 a UE can only be transmitted to on a single spatial layer and only transmission to orthogonal UE pairs can be performed, meaning that none of these techniques could yet be used. With the introduction of non-codebook based precoding (enabled by the DM-RS) in LTE-A Rel.10, however, support for ZF and BD was enabled. However, due to the presence of quantisation errors in the CSI feedback, use of MU-MIMO often requires interference suppression at the receive side, which requires multiple Rx antennas [25]. When BD is used the combination of UE antenna requirements for multiple spatial streams *and* interference suppression results in the number of UE Rx antennas needed being quite high. For this reason we focus on Zero-Forcing Beam Forming for MU-MIMO operation in this thesis, implementation of which is discussed later in Section 3.2.3, also the equations in the previous subsection are based on ZFBF transmissions.

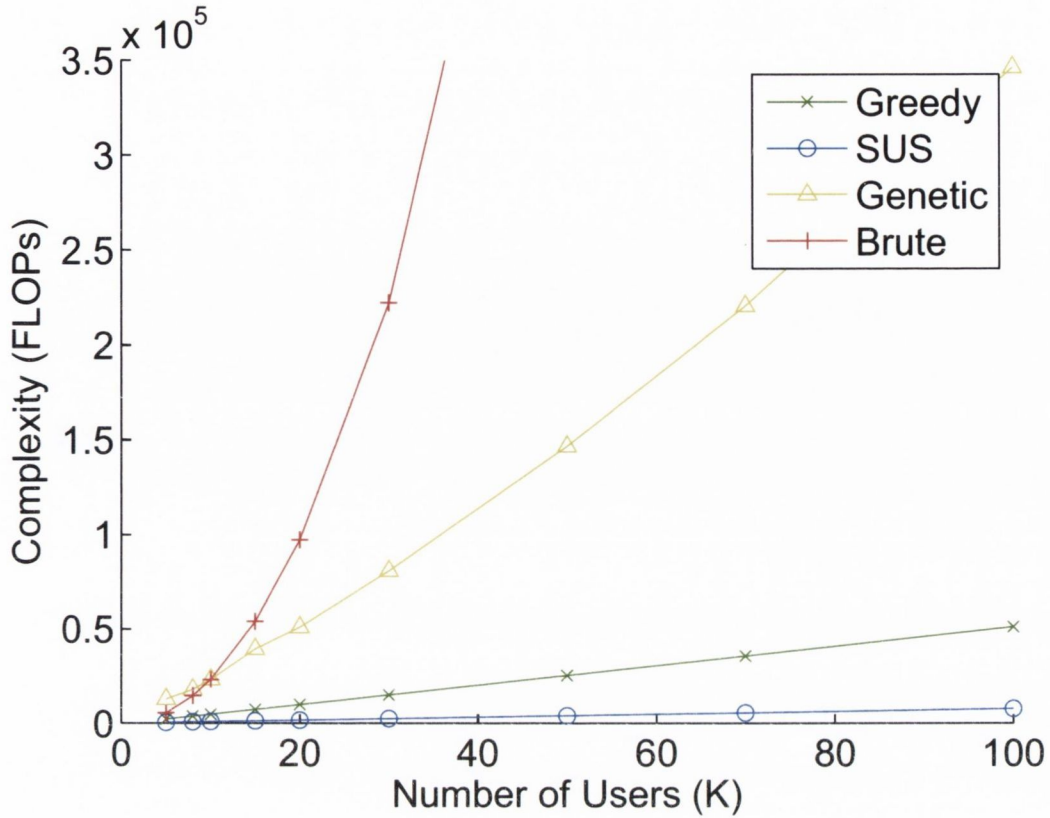


Figure 2.9: ZFBF 2×1 computational complexity comparison of MU-MIMO schedulers.

MU-MIMO Scheduling

User scheduling is the most common method used to take advantage of multi-user diversity. However, in MU-MIMO systems, user scheduling becomes a lot more complex than in SU-MIMO systems. This is because the system performance is not just a function of which UE is selected, but also of which UEs it is paired with; for example, while in SU-MIMO taking a brute force approach would require computing the objective function K times (where K is the number of UEs in the cell), trying all possible MU-MIMO UE combinations to find the optimal set of n_{MU} UEs requires $\sum_{n_{MU}=1}^{N_{MU}} \binom{K}{n_{MU}}$ objective function computations, where n_{MU} is the number of UEs co-scheduled for MU-MIMO with a maximum of N_{MU} . As a result greedy heuristics are commonly used.

In figures 2.9 and 2.10 we compare the complexity and performance of a number of different scheduling algorithms for 2×1 ZFBF MU-MIMO. The algorithms considered are: the Greedy User Selection (GUS) algorithm proposed in [26] which

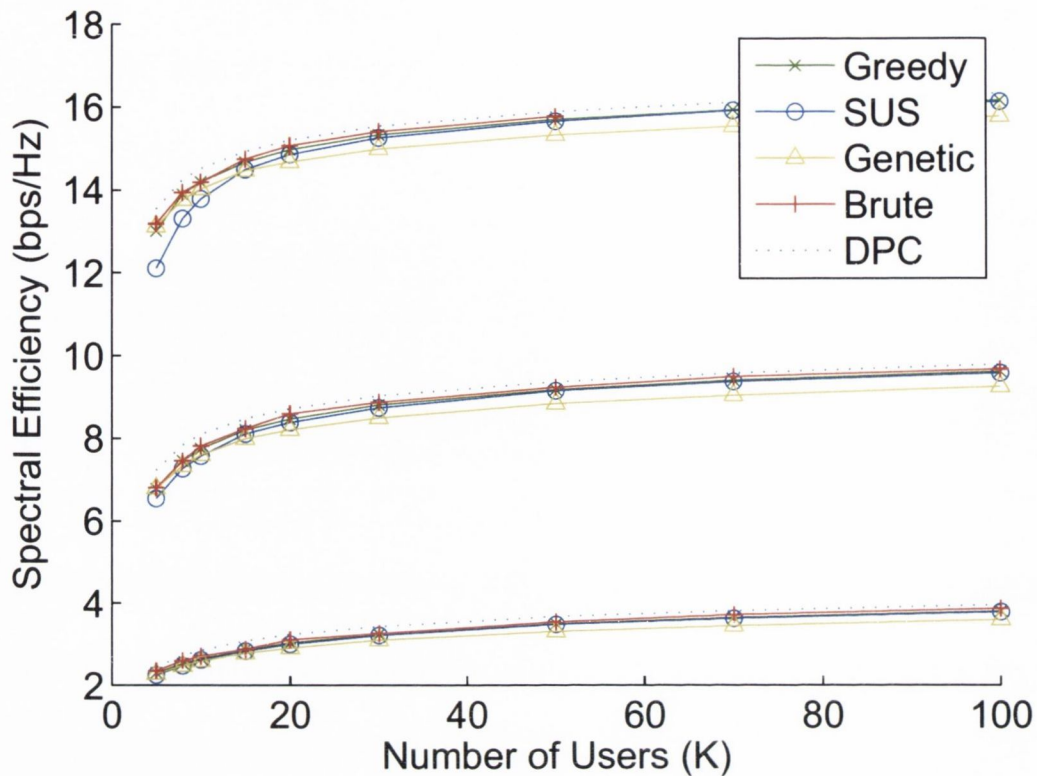


Figure 2.10: ZFBF 2×1 spectral efficiency comparison of MU-MIMO scheduling algorithms. From bottom to top the three bands of curves correspond to SNRs of 0dB, 10dB and 20dB, respectively

successively selects UEs based on their expected rates (*Greedy*); the Semi-orthogonal User Selection (SUS) algorithm in [27] which successively selects UEs based on their channel gain, provided that their fed back precoding vector is semi-orthogonal to those of the already scheduled UEs (*SUS*); the heuristic genetic algorithm proposed in [28] (*Genetic*); and the brute force approach. For comparison Figure 2.10 also includes the DPC performance upper bound.

These figures were generated using Matlab-based code made available online at [29] as part of the publication [30] to which we added the investigated scheduling algorithms. Flop counting for complexity comparison was computed using the Lightspeed Matlab toolbox [31]. We performed these simulations at link level assuming perfect channel knowledge at the transmitter, with the maximum number of co-scheduled UEs $N_{MU} = N_t = 2$, and for a range of SNR values as shown in Figure 2.10 [32].

As can be seen, the rapid growth in the brute force complexity means that it

would be impractical in real systems. On the other hand, the SUS algorithm, which we use primarily in the following chapters, provides an extremely low complexity solution which also achieves close to optimal performance for high numbers of UEs, and while the spectral efficiency is reduced when the number of UEs per cell is low, overall this algorithm achieves a good tradeoff between complexity and performance¹.

2.1.3 Coordinated Multi-Point (CoMP)

In this thesis our main focus is on the gains of downlink multi-antenna techniques specifically focussing on small cell network scenarios with coordination. This focus is shared with the more generalised group of techniques known as Coordinated Multi-Point (CoMP). From a theoretical perspective the potential gains of CoMP have been known for a considerable amount of time [33]; however, the costs involved in implementing the required coordination have, until recently, outweighed the benefits. The recent uptake of new network architectures which centralise the processing of multiple eNBs such as Cloud Radio Access Network (C-RAN) [34] and liquid radio/baseband pooling [35] have led to cost effective solutions which greatly simplify and accelerate coordination between neighbouring small cells. Amplifying this, the emphasis of the 3GPP on supporting CoMP functionalities in LTE-A systems (first discussed for Release 9 [36] with significant steps made in Release 11 [37]) has meant that this is a hot topic at the moment which is generating a lot of interest from the research community [38]. Further emphasis on this sort of investigation has been encouraged by discussions such as [39], which asks questions such as "Is the PHY Layer Dead?" provoking a shift in telecommunications research from a backdrop of incremental link level improvements toward new network approaches with potentially high undiscovered gains.

The idea behind downlink CoMP is that UEs on the cell edge can often receive transmissions from multiple base stations at once; if these transmissions are coordinated then the performance of these UEs can be considerably increased. Very good resources on the CoMP background and literature can be found in [38, 40, 41], as well as the 3GPP document on CoMP itself [37].

CoMP operates in three main ways: Coordinated Scheduling/Coordinated Beamforming, Dynamic Cell(/Point) Selection, and Joint Transmission. These

¹Even at $K=4$ the brute force algorithm requires 14.6 times as many computations as SUS.

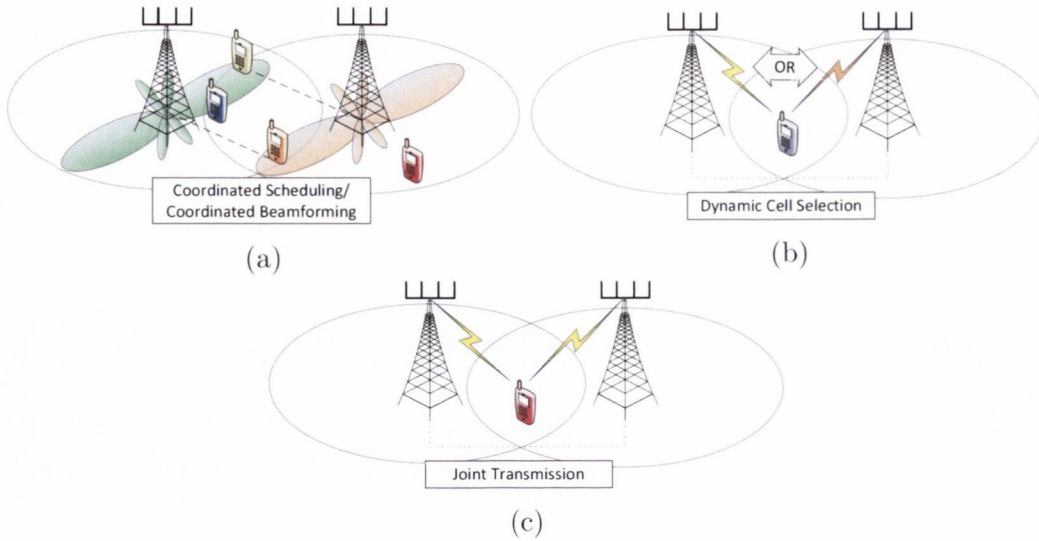


Figure 2.11: CoMP transmission Types. (a) Coordinated Scheduling/Coordinated Beamforming, (b) Dynamic Cell Selection, (c) Joint Transmission.

transmission types are depicted in Figure 2.11. In this subsection we focus on providing background on these three CoMP types, while the relationship between these and our work will be discussed more precisely in Chapters 4 and 5.

Coordinated Scheduling/Coordinated Beamforming

In Coordinated Scheduling/Coordinated Beamforming (CS/CB) transmissions to each UE originate from a single serving base station(/transmission point) and the scheduling processes of the set/cluster of cooperating eNBs (CoMP set) are coordinated. As shown in Figure 2.11a, CS/CB operates by scheduling UEs such that the intercell interference (marked in dashed lines) from the beams of the scheduled UEs toward UEs operating within the same time/frequency resource in adjacent cells is minimised. This is mainly done in two ways, which are referred to in [40] as Coordinated Beam-Switching (CBS-CoMP) and Coordinated Scheduling (CS-CoMP).

In *CBS-CoMP* base stations specify certain resources in which different beams can be used (beam-to-resource assignment). This can be done in a time-domain manner in which beams are cycled between according to a pattern specified through coordination [42, 43], or in the frequency domain in which groups of subbands are assigned to each beam [44, 45], then, within each subband group, only UEs using the

corresponding beam can be scheduled. CBS-CoMP provides effective interference reductions at a low coordination signalling cost; however, these distributed solutions are only suitable in cases where the number of UEs per cell is large, unlike the scenarios we consider.

In *CS-CoMP* methods a different approach is applied in which cell-edge UEs inform their serving eNB of the set of either their worst-interfering (or least-interfering) beams from neighbouring cells. This information is then fed to the neighbouring cells which either avoid (or steer toward) usage of those beams in the scheduling decisions [46], [38, Sec 5.3 and references therein]. These methods are largely distributed and require inter-cell feedback exchange in the form of beam indexes (e.g. PMIs for the LTE Rel.8 predefined codebook of beams) and the corresponding channel quality improvement (CQI improvement for LTE).

As specified in [37, Sec 5.1.3] CS/CB can also be performed in tandem with Semi-Static Point Selection (SSPS) in which transmissions to a UE originate from a single transmission point(/eNB) which can change semi-statically (on the order of seconds). SSPS as a term appears to be as-of-yet relatively unused in the research literature, although is very closely related to its more dynamic counterpart, DCS, discussed in the following.

Dynamic Cell Selection

Dynamic Cell(/Point) Selection (DCS) is a form of CoMP Joint Processing (JP), meaning that transmit data must be present at all eNBs within the cooperating CoMP set. Despite this, transmissions to the served UE originate from a single transmission point, as shown in Figure 2.11b; although the transmission point can be dynamically switched from one subframe ($1ms$) to the next. DCS is often performed in along with *muting* [47] in which the transmissions of cooperating eNBs, when not transmitting to the dynamically switching cell-edge UE on a given subband, are halted temporarily (replaced with Almost Blank Subframes (ABSs)), meaning that their interference to the cell-edge UE is removed.

Joint Transmission

Joint Transmission (JT) is the simultaneous transmission of data from multiple transmission points to either a single UE, or multiple UEs as shown in Figure 2.11c. As JT requires the transmit data to be present at all eNBs within the CoMP set, this is the second CoMP downlink JP technique.

Essentially JT CoMP combines the channel matrices of multiple cells into a single channel matrix to which MIMO transmission techniques are applied. This can be done *coherently* or *non-coherently*, which are analogous to *closed-loop* and *open-loop* operation in single-cell MIMO, as were discussed in Section 2.1.1. As an example of non-coherent CoMP JT, [48] considers the use of (open-loop) Space-Frequency Block Coding to achieve Transmit Diversity (TxD) gains for cell-edge UEs, which requires a lower feedback overhead than coherent CoMP solutions.

While coherent CoMP JT systems require significantly more feedback overhead than other CoMP schemes, it is these techniques that hold the highest potential gains. These techniques operate by performing (often multi-spatial-layer) precoding across multiple cells at once, essentially turning the CoMP set into a single distributed antenna array. This can focus on serving either a single cell-edge UE, for which the signal quality will be dramatically increased, or multiple UEs simultaneously in a similar way to MU-MIMO [38, Sec 6.3, 6.4], for which, under full perfect channel knowledge with Dirty Paper Coding (DPC), the performance is lower bounded by that of a single isolated cell [49] (corresponding to the removal of all adverse interference effects). This multi-user CoMP method is also commonly known as *Network MIMO* or *Virtualised MIMO* and has met with a lot of interest in the research community over the years [50]. It is also the main focus of downlink C-RAN architectures which aim to provide a feasible means of implementing these systems in reality. As a result, a number of C-RAN-centric works also provide comprehensive CoMP JT solutions [51]. This and other architectures for coordination between small cell nodes will be discussed in more detail in Section 2.2.2.

CoMP Channel Feedback

In CoMP feedback two important considerations exist: what to feed back, and how to obtain that information. A discussion of these considerations focusing on CoMP




Figure 2.12: LTE-A Downlink Reference Signal Structure for 2 Tx antenna configuration.

JT can be found in [38, Sec 9.1].

In Rel.11 of LTE-A two general types of CoMP channel state feedback for Frequency Division Duplexing (FDD) scenarios are discussed [37, Section 5.2.2]: *explicit* and *implicit* channel state/statistical information feedback. Explicit feedback consists of quantised channel, noise and interference matrices for each cooperating eNB (point), which can amount to quite a high feedback overhead and as such should only be considered for JT scenarios. The reduced overhead alternative is implicit feedback, which uses similar feedback to single-cell LTE scenarios (Channel Quality Indicator (CQI), Precoding Matrix Indicator (PMI), Rank Indicator (RI)) although this is potentially fed back for each cooperating point, and subject to a number of hypotheses regarding Tx/Rx processing. These hypotheses can range from SU-MIMO or MU-MIMO usage, to whether single or multi-point transmissions are to be used, to assumptions on the interference beamforming used in different sub-bands or time-slot in cases where CBS-CoMP is used.

To avoid the need to introduce new reference signalling types, CoMP CSI is obtained from the same reference signalling as single-cell CSI. In LTE-A systems

this consists of three types: Cell Specific Reference Signal (CS-RS) (a.k.a. Common Reference Signal (CRS)), CSI Reference Signal (CSI-RS), and Demodulation Reference Signal (DM-RS) (a.k.a. UE-specific Reference Signal (URS)). Figure 2.12 shows an example reference signal structure for LTE-A for a single Resource Block (RB) with 2 Tx antennas(/antenna ports).

In LTE Rel. 8 systems the CS-RS was the only reference signal present and was used for both CSI estimation and demodulation. CS-RSs consists of non-coded reference signals for each antenna port, for which many are needed.

In LTE-A systems, which must support transmissions from up to 8 antenna ports at once, it was no longer feasible to only use the bulky CS-RS and instead the operations of CSI estimation and demodulation were separated to be performed using the CSI-RS and DM-RS, respectively, instead [52][53, Sec 6.4]. As CSI estimation doesn't require as much reference signalling as demodulation the CSI-RS can be quite sparse, requiring only one Resource Element (RE) per antenna port as shown in Figure 2.12. Unlike the CS-RS and CSI-RS, the DM-RS is precoded in each RB in the same way as the REs used for data transmissions, which simplifies the demodulation process considerably. As the precoding used on the data channel of a given RB is specified based on the UE that is scheduled within it, this is a UE-specific reference signal and as such the required number of DM-RSs can scale with the number of spatial layers used by the scheduled UE, rather than being fixed at the maximum possible (the number of antenna ports) as was the case for the CS-RS. This corresponds to a much reduced reference signalling overhead in cases where there are many Tx antennas and a small number of spatial layers used.

Nonetheless, for backward compatibility reasons and for a number of other functionalities such as cell association and handover, the CS-RS is still required in LTE-A systems. This is fortunate, as for CoMP scenarios, the CS-RS offers the primary means of neighbouring cell CSI estimation. This is because the locations of the CS-RSs in neighbouring cells are shifted in the frequency domain based on a mod (6) of their cell ID. This corresponds to circularly shifting all of the CS-RS locations upwards in Figure 2.12. This was originally intended to provide frequency-domain orthogonality between the reference signals of neighbouring cells in a macrocell hexagonal grid structure, although the CS-RSs are also coded with cell-specific Zadoff-Chu sequences to provide orthogonality in most other cases.

In the case of the CSI-RS, the abilities to specify different CSI-RS RE locations

in adjacent cell and for base stations to mute transmissions in a given CSI-RS RE location to allow for neighbouring cell readings (known as zero-power CSI-RS transmission), combined with the use of cell-specific Hadamard sequences, ensure that neighbouring cell CSI estimations using the CSI-RS are supported.

In order to support the transport of this increased amount of feedback from the UE to the serving base station a new control channel in LTE-A of adjustable size called the Enhanced-PDCCH has been defined [54]. Sharing of information between neighbouring base stations is performed over what is called an X2 connection, and the speed of this connection depends on the connection type (e.g. high speed fibre) and the network architecture; for example, information exchange between centrally processed base stations can be assumed to be instantaneous.

In our work, UEs will be required to compute implicit PMI and CQI feedback for neighbouring small cells. The hypotheses associated with this feedback will change for each of the two investigated coordinated reassignment scenarios in Chapters 4 and 5.

2.2 Advances in Small cells

To cope with ever increasing traffic demands, Mobile Network Operators (MNOs) are constantly looking for solutions to enhance the capacity of their network while keeping deployment and operational costs to a minimum. By significantly improving the spatial reusability of spectrum resources, small cells have shown huge potential to address the capacity deficit that mobile operators are expected to face in the near future. Additionally, the low cost of small cell deployment and operation, compared to macrocell base stations makes them a very appealing technology [55]. As a result the number of small cell deployments worldwide is now over 8.4 million [3], deployed by more than 64 operators. The vast majority of these are currently residential femtocells (8.1 million) which were the first to be widely deployed; however, other deployment types are gaining pace fast as standardisation for each field is introduced [56]. These include: enterprise small cells, indoor urban small cells and outdoor urban small cells, and rural small cells.

In this section we focus on three aspects of small cells which are of importance in this work: multi-antenna usage, coordination, and energy savings.

2.2.1 The Use of Multiple Antennas in Small Cells

While a large proportion of small cell research focuses on single antenna transmissions, given the clear benefits of multi-antenna downlink transmissions in macro-cellular networks it makes sense to consider multi-antenna techniques in small cell networks too [57, Ch 5].

Given the interference limited nature of small cell scenarios, multi-antenna transmissions in this area tend to focus on beamforming solutions as they increase served UEs' SINRs while limiting interference to UEs in neighbouring (macro/small) cells. Further, given the small form factor of small cell base stations (hence antennas needing to be closely spaced) and the high likelihood of strong channel correlations in small cell environments (resulting from strong LOS components and reduced scattering compared to macrocell cases [5]), these scenarios are particularly suited to beamforming transmissions. These are mainly performed in two ways: using switched Multi-Element Antenna (MEA) systems [5, 58, 59], and by applying linear precoding to active antenna arrays [60–63].

The use of switched MEA systems is a beam-steering solution designed by Bell Laboratories, Alcatel-Lucent for low-cost small cell deployments. In this, multiple antennas, each of different antenna gain patterns (e.g. patch antennas – which essentially form a beam in a single 90° sector), point in different directions. In [58] these MEAs are used to confine Closed Subscriber Group (CSG) small cell coverage to the interior of a home so as to focus transmissions on indoor coverage areas and limit leakage to external macrocell UEs. Use of MEAs is also investigated in [5] for mobile network operator-deployed picocells in an urban environment (Munich city centre, see Sec 3.3.2) to improve UE SINRs and reduce the small cell density required to meet demand in the investigated urban area.

The ability to adaptively refine the coverage area of multi-antenna-equipped small cells is particularly useful in urban small cell deployments as it is not always possible (due to site acquisition limitations, etc.) to place a small cell in its optimal location. In these cases the coverage area can be adapted to serve the desired area and can also be reconfigured based on changing traffic dynamics (for example between day and night).

Of the scenarios investigating the application of linear precoding to active antenna

arrays in small cell scenarios, [61] and [62] are of particular interest to us due to their investigation of MU-MIMO transmissions.

In [61] the authors investigate the coverage and spatial reuse gains that can be achieved from SU-MIMO and MU-MIMO with ZFBF and no user scheduling, in two-tiered networks subject to a minimum SINR constraint for successful transmission and assuming perfect channel knowledge. In small cell networks, increased spatial reuse can be achieved in two ways: through increasing the small cell density or through use of MU-MIMO. When ZFBF MU-MIMO is used there is a reduction in SINR as some of the power that would be used to boost signal strength in single-layer transmissions is instead used to reduce multi-user interference through zero-forcing. This means that, for a given minimum SINR threshold, small cells utilizing MU-MIMO transmissions cannot be deployed as densely as those performing SU-MIMO. As a result the authors find that the spatial reuse gains from using a more dense SU-MIMO deployment exceed those from MU-MIMO usage. However, this work neglects to consider user scheduling or the use of a semi-orthogonality constraint, both of which lead to higher orthogonality between co-scheduled MU-MIMO UEs and in doing so decrease significantly the reductions in SINR which result from the ZFBF operation.

In [62] the work of [61] is extended to also consider user selection (scheduling). The authors show that, in this case, MU-MIMO performance exceeds SU-MIMO performance in small cell networks, although the number of UEs to serve simultaneously, n_{MU} , depends on a number of factors such as the deployment density, the number of UEs present in each small cell (K), the pathloss exponent, and the minimum SINR threshold. For this reason an adaptive technique is proposed which switches n_{MU} to achieve the best performance in terms of the sum of the expected rates of the UEs served within the small cell taking into account the ZFBF operation. To summarise the conclusions, it is found that the optimal n_{MU} increases as the deployment density decreases and as the number of UEs per cell increases (for which a range of $K=4$ to $K=16$ is considered). The n_{MU} -adapting user selection scheme applies a brute force approach in which all possible UE combinations for all values of n_{MU} are considered. As a result, even though no semi-orthogonality constraint is applied, the best UE groupings, which in cases of many UEs per cell will be of high SINR and high mutual channel orthogonality, will be selected.

In this thesis, we specifically investigate cases where we are constrained by low numbers of UEs per small cell. In these, we consider the reassignment of UEs between neighbouring small cells to reach a configuration in which the channel orthogonality between UEs within a given small cell is increased. In doing so we reduce the losses in SINR from ZFBF and increase the number of UEs for which MU-MIMO performance exceeds SU-MIMO in the given small cell scenario.

2.2.2 Collaboration and Supporting Network Structures

Although it ventures slightly outside of the scope of this thesis, it is worthwhile to be aware of different network architectures which have the potential to enable coordination between neighbouring small cells, where necessary. To provide an overview of what we consider later on in this thesis: we investigate a network architecture in which small cell base stations coordinate with a centralised controller which issues reassignment and sleep state usage commands. As we will consider a low mobility scenario (slow changing channel) in which coherent joint transmissions are not performed (UEs are served by a single eNB at a time) tight synchronisation between cells and fast precise coordinated information exchange will not be necessary, meaning that a number of suitable architectures which could provide such coordination exist. These are discussed within this subsection.

Within LTE networks coordination between eNBs is performed over two interfaces, namely, the X2 interface, which is a logical connection between all neighbouring eNBs, and the S1 interface, which connects the eNB to the Mobility Management Entity (MME) and the Serving GateWay (S-GW). In the traditional view of coordinated network architectures, all transceiver processing and user scheduling is performed by each eNB in a distributed manner and only required signalling is shared among cooperating nodes to provide sufficient functionality for our desired operation. The centralised controller can then be located either at one of the cooperating eNBs, at the MME, or as a separate entity. In cases where the frequency of the reassignments is higher (i.e. they occur more often, as for Dynamic Cell(/Point) Selection), within this architecture reassignments can be sped up if the data to transmit is simultaneously buffered at multiple cooperating eNBs so that the data is already in place should a reassignment command be issued. For our investigated small cell scenario this mostly-distributed architecture could provide sufficient coordination capability.

Another solution to the coordination of cellular networks, which is currently growing in popularity, is to take advantage of deployment architectures which perform the baseband processing of multiple small cell eNBs centrally at a single location. In order to allow fast flexible deployment of nodes for next-generation rollout or in cases of increased traffic demands, as well to simplify network software updates and the implementation of Self-Organising Networks (SON) functionalities, there has been a lot of interest from industry in fronthauling architectures [64]. In these, the small cells take the form of simple Remote Radio Heads (RRHs) from which the received radio signals are carried to and from a centralised BaseBand Unit (BBU) which performs the baseband processing, scheduling and higher layer functionalities all at a single location. Having the baseband processing of multiple base stations at a single location greatly facilitates coordinated operation as extremely fast information exchange can be performed between co-located cooperating base stations. This architecture makes it possible to support joint decoding and joint processing/precoding of signals from multiple base stations at once as in JT CoMP and network/virtual MIMO [50] and is the main mode of operation for C-RAN networks [34, 51, 65]. This fronthauling concept has been given a plethora of different names by different companies/research groups including light radio [66], radio over fiber [67, 68], liquid radio/baseband pooling [35], Radio Dot Systems [69], and cloud-based architectures. However, as these architectures are still within a developmental phase current deployments offer only simple coordination while future deployments will offer more advanced capabilities.

Another interesting architecture for heterogeneous networks which is growing rapidly in popularity is that of Phantom cells [70]. In this the heterogeneous deployment concept is combined with the carrier aggregation concept, leading to a scenario in which the macrocell base stations operate in a low frequency band and perform all control-plane operations, while the small(/phantom) cells operate in a higher frequency band and are only required to send UE-specific signals. Beyond dramatically reducing the control/reference signalling overhead in the phantom cell transmissions this opens up a range of possibilities for coverage/handover/cell reassignment, coordination, massive MIMO and mmWave applications. This architecture creates a scenario in which the small cell connection/assignment becomes transparent to the UE as the control plane is provided by the covering macrocell which dramatically decreases overheads involved in reassigning UEs

between cells. As a practical consideration this concept is envisioned to also make use of RRHs, in which case the signal processing of all small cells within the coverage area of a macrocell base station will be performed centrally at the macrocell site.

2.2.3 Green Wireless

As well as enabling large gains in spatial reuse, network densification through use of small cell technology has the potential to hugely reduce the cellular industry carbon footprint [71]. As Radio Frequency (RF) transmit powers roughly scale with the transmission distance to the power of the pathloss exponent, $\alpha > 2$, it is easy to see how the deployment of small cells, by reducing the distance between the base station and the UE, can lead to a reduction in required transmit powers. This is even more pronounced for indoor UEs where, by placing small cells eNBs indoors, external wall penetration losses can also be removed and the pathloss exponent often lowered. Further reductions still can be achieved as the non-Tx power consumption of a small cell base station device (e.g. baseband processing, cooling, etc.) is significantly lower than that of a macrocell base station.

That said, given the massive scale of expected small cell deployments it is predicted that small cell technologies will have the effect of adding an additional 4.4 terawatt-hours(TWh) or 5% to the total global power consumption of cellular networks by 2020 [72]. For this reason a number of ways to reduce small cell power consumption further have been proposed; these include: tailoring the small cell deployment toward maximum energy efficiency (e.g. the Cell-on-Edge concept in which small cells are around the cell edge rather than uniformly with the macrocell coverage area [72]), use of base station sleep modes to reduce power consumption when base stations are inactive, cell breathing in which small cells adapt their cell size based on dynamic network load, equipping small cell base stations with renewable energy sources (e.g. solar panels), power amplifier improvements, and use of relays. In this thesis, we focus primarily on energy savings achieved through the use of dynamic sleep modes.

In order to quantify and assess power consumption and potential energy savings within cellular networks the European Commission 7th Framework Programme (FP7) project Energy Aware Radio and neTwork

tecHnologies (EARTH) [73] has put together a framework known as the Energy Efficiency Evaluation Framework (E³F) which takes into account base station power models and traffic demands to estimate both short-term and long-term total network power consumption. The power model employed takes into account sleep mode usage (small cell eNBs can operate in either an *active* or *idle* state) and is expressed as follows [74, 75]:

$$P_{in} = \begin{cases} N_t(P_0 + \Delta_P P_{out}), & ACTIVE \\ N_t P_{sleep}, & IDLE \end{cases} \quad (2.7)$$

where P_0 represents the power consumption at zero RF output power, P_{out} represents the RF output power which has a maximum of P_{max} and Δ_P is the slope of the load dependent power consumption. For full buffer traffic, as used in this work, transmissions will operate at maximum power, meaning that P_{out} for active eNBs will be equal to P_{max} . P_{sleep} represents the power consumed by an eNB in sleep mode. An eNB can switch to sleep mode to save energy if there are no UEs requesting transmissions which it has to serve. These power consumption values will differ depending on the type of base station considered, be it femto, pico or macro.

We will use this power consumption model later on in Chapter 5. Related work regarding usage of small cell sleep states will also be discussed in Chapter 5.

2.3 Summary

In this chapter we have provided the reader with an overview of the background and state of the art, relevant to this thesis, in the areas of multi-antenna transmissions and small cell networks.

We started by reviewing SU-MIMO transmission types and the cases to which each is best suited, before changing our focus to MU-MIMO transmissions in which multiple UEs are simultaneously served on distinct spatial layers within a single time/frequency resource. We looked at a number of different MU-MIMO transmission strategies and scheduling approaches with a primary aim to provide necessary background before we proceed onto the discussion of our MU-MIMO simulation implementation in the following chapter. This was followed by a

discussion of the different CoMP transmission types, the ways in which they coordinate their transmissions, the feedback exchange that this requires and how CSI for neighbouring cells can be effectively acquired.

Our discussion on advances in small cells centred on three main areas: the use of multiple antenna transmission in small cell networks, supporting architectures for small cell coordination and energy savings in small cell networks.

The following chapter will focus on the simulation of our investigated scenarios and will build on the background provided in this chapter toward the assessment of our proposed MU-MIMO-based reassignment mechanism in chapters 4 and 5.

3 Simulation Methodology and Scenarios

In this thesis our objective is to investigate the potential benefits of enabling *Multi-User MIMO (MU-MIMO) Across Small Cells*. This investigation relies, at least partly, on simulation. For this, we require the ability to simulate heterogeneous network scenarios in which inter-cell interference, multi-antenna capabilities and User Equipment (UE) scheduling are accurately modelled. The simulations must support coordination, information exchange and UE reassignment between adjacent base stations. Further, these simulations must be performable within a reasonable timescale.

In this chapter, we conduct an assessment of potential simulation tools to be used in our study. In the process, we provide a comparative analysis of simulators available for Extremely Dense Networks (EDNs) in 4G and beyond. This is broader than simply the selection of a simulator for our study, and provides guidelines for the selection of an appropriate simulator in future studies of EDNs.

We identify the missing components of the simulator to enable our study of *MU-MIMO Across Small Cells*, and design and implement those components. We plan to make these extensions available to the broader community as they will be useful in future studies of multi-antenna capabilities in dense small cell deployment scenarios.

3.1 Comparison of Simulators

The accepted definition of Heterogeneous Networks (HetNet) has gradually evolved toward extreme network densification. Researchers, industry, and network operators are counting on the massive deployment of small cells as a key coping strategy for the foreseen data tsunami. The high throughput gain offered by the addition of these low power base stations is hoped to provide a solution that is efficient in terms of energy, spectrum and cost. Nonetheless, technological advances always go hand in hand with new technical challenges.

As illustrated in Figure 3.1, Extremely Dense Networks (EDNs) are characterized by massive small cell deployments which enable the operator to offload low-mobility user traffic away from the macrocell network. Despite the clear advantages of these networks, their high density can potentially result in high inter-cell interference, bottlenecks in the backhaul and increased energy consumption. To address these challenges, international and European research projects, like CROWD [76] and METIS [77], have begun exploring flexible network architecture designs and user-cell association procedures.

Due to the scale of these scenarios, simulation forms an essential tool in investigating potential EDN solutions. Here, discuss what problems one might face when simulating these new dense scenarios. The main contributions of this section are in the identification, outlining and addressing of key simulation considerations for EDNs. These include, firstly, choosing the best starting point, whether it be commencing from scratch or extending existing simulators (and if so, which ones?) taking into account the investigation requirements; secondly, identifying supported features, e.g. implemented backhaul protocols or support for interference suppressing receivers; and thirdly, taking into account practical considerations such as ease-of-use, extensibility and run time.

We compare a representative set of potential EDN simulators looking at popular openly available system and network level simulators, as well as including an in-house-built system level simulator with reduced functionality but more specific EDN focus.




Figure 3.1: SINR distribution in an extremely dense heterogeneous deployment.

3.1.1 Simulator Types

To introduce the various approaches used to simulate EDNs, this section differentiates simulators according to the layers of the protocol stack on which they focus, shown in Figure 3.2. Representative simulators are also identified for later comparison.

Focus

Simulators are generally categorised as either link, system or network level, which is reflected in their implementation.

Link Level (LL) simulation tends to focus on the performance of Physical (PHY) properties, usually for a single link between a base station and a user, either in the uplink or downlink. To obtain accurate single-link Block Error Rate (BLER) statistics, LL simulators usually implement all blocks of the PHY in detail with minimal abstraction. In the EDN context, LL simulators can be used to validate the

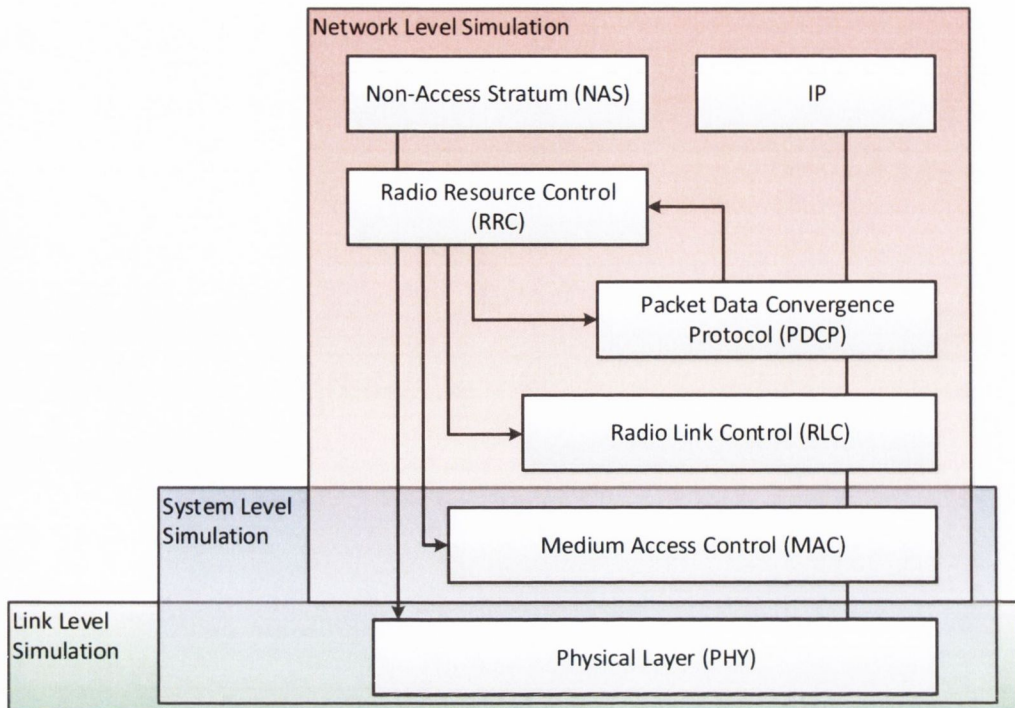


Figure 3.2: Mapping between simulator types and LTE protocol stack layers.

performance of system and network level simulators, which rely on BLER data from LL simulation to abstract the PHY.

When studying complex systems consisting of several base stations and user terminals, **System Level (SL) simulators** are preferred to LL due to their reduced simulation times. SL simulators enable assessment of coverage, spectral efficiency, and throughput in wireless networks, taking into account interference generated by neighbouring evolved Node Base stations (eNBs) or UEs. SL simulators are also suitable for studying how Radio Resource Management (RRM) and interference coordination techniques perform and how potential solutions scale with the size of the network.

Unlike SL simulators, which tend to be protocol agnostic and focus on the air interface, **Network Level (NL) simulators** are designed to facilitate investigations into specific protocols and their interactions with the upper and lower layers. Moreover, by treating base stations as network entities capable of exchanging messages among themselves, NL simulators can also model and assess EDN backhaul issues.

Representative simulators

For this section, we chose three simulators to compare, assessing which are best suited to different aspects of EDNs. These three simulators are the *Vienna SL simulator* v1.8r1375 for Long Term Evolution (LTE) downlink [8], the NL *ns-3* LTE module [78] v3.21 and an in-house-built custom SL simulator, which we will refer to as *HetDenSim*; all three simulators are designed for LTE, which we consider to be a good baseline for EDN simulations. No LL simulator was selected, as EDNs focus on large network – rather than single link – interactions.

HetDenSim targets Heterogeneous Networks and large networks of small cells. It was designed with time and frequency domain based RRM techniques in mind and is particularly suited to coverage, Signal to Interference and Noise Ratio (SINR) statistic, and expected throughput investigations in dense networks. To speed up large network simulations, HetDenSim uses a simplified physical layer model. This comes at the cost of simulator limitations concerning algorithms and techniques that require detailed PHY abstraction. HetDenSim’s populous small cell network design emphasis makes it a well-suited tool for EDN simulations.

The Vienna SL simulator, which is freely licensed for academic use, has been designed to assess both Homogeneous and Heterogeneous LTE network deployments. The Vienna SL simulator implements the PHY with a low level of abstraction and therefore, from the EDN perspective, provides a powerful means of testing techniques which require detailed PHY implementation, such as Multiple Input Multiple Output (MIMO).

ns-3 is an open source general purpose NL simulator for both wired and wireless networks; in particular we will mainly refer to the LTE module for *ns-3* [78]. Given its extensible nature, *ns-3* could provide a viable EDN-specific simulation environment.

With these three simulators we aim to provide a cross-section of the generalised simulator types best suited to EDN simulation. Vienna SL and *ns-3* representing widely accessible SL and NL simulators, respectively, and HetDenSim representing purpose-built custom simulator solutions.

3.1.2 Supported Features

In this subsection we outline some important features/functionality and associated implementation methods for Extremely Dense Network (EDN) simulation.

Link Level Abstraction

Full implementation of all links at bit-level would impose a heavy computational burden on system level or network level simulators, especially when the number of links to simulate grows large. To deal with this problem, link level abstraction (also known as a Link-to-System (L2S) interfacing) is used [38, Section 14.1]. In this, subband SINRs are mapped(/averaged) to a single effective SINR, γ_{eff} , from which the expected BLER can be computed. The mapping to the effective SINR can be expressed as

$$\gamma_{eff} = \frac{1}{\beta} I^{-1} \left(\frac{1}{B} \sum_{b=1}^B I \left(\frac{\gamma_b}{\beta} \right) \right) \quad (3.1)$$

where $b \in \{1, \dots, B\}$ are the subbands for which γ_b is the subband SINR, β is a calibration factor obtained from an L2S training process and I is the non-linear mapping function which is determined by the Effective SINR Mapping (ESM) method used (e.g. Exponential ESM (EESM), Logarithmic ESM (LESM), Mutual Information ESM (MIESM), etc.). MIESM is commonly used in L2S interfacing because it is known to have better performance than most other techniques [79]. Both the Vienna SL simulator and ns-3 perform MIESM.

To map the effective SINR to the corresponding BLER, a set of curves (one corresponding to each Modulation and Coding Scheme (MCS)) are obtained from LL simulations, for example using the Vienna LL simulator [16]. By way of illustration, these curves are plotted in Figure 3.3; the MCSs in this case correspond to Channel Quality Indicators (CQIs) 1 to 15.

Once obtained, the BLER, combined with knowledge of the used MCS and bandwidth (for which combination the corresponding number of transmitted bits will be known) enables the determination of the average number of successfully decoded bits within each transmitted frame.

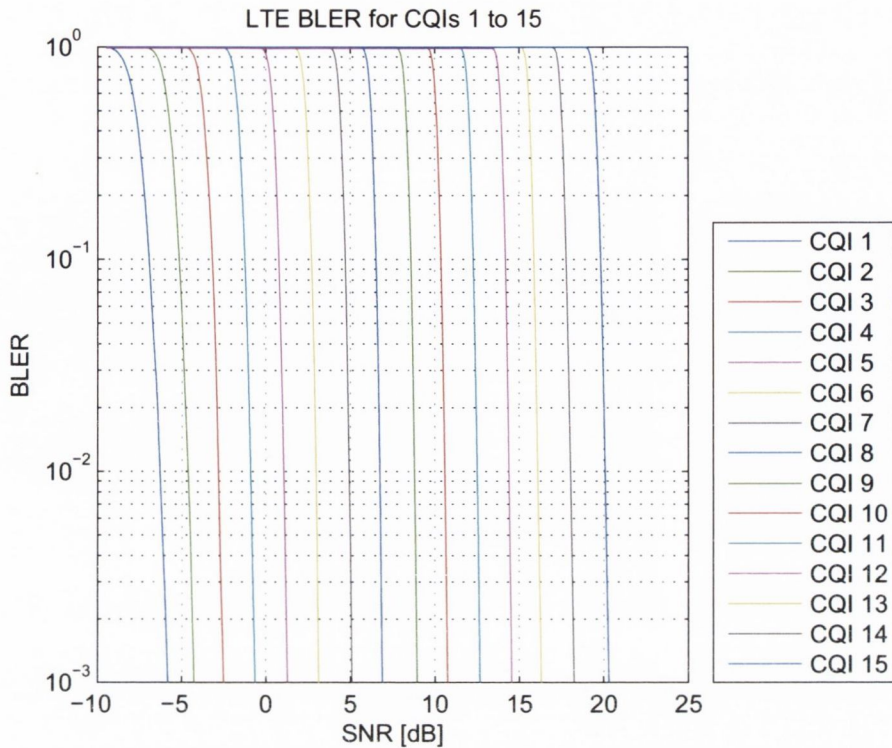


Figure 3.3: Curves used for mapping Effective SINR to BLER for a given CQI.

Less computationally intensive methods such as the modified Shannon capacity formula [80] also exist for computing UEs' data rates in LTE systems. HetDenSim adopts this approach for the PHY. These methods often lack flexibility, especially when taking into account new features, feedback uncertainty or changing network dynamics.

The L2S interface is of particular importance to EDNs if one needs to implement interference cancellation or multi-antenna transmission techniques. As these techniques work at the PHY level, detailed abstraction of the PHY provides more accurate simulation results.

Interference Suppressing Receivers

As the network topology develops toward a more numerous, heterogeneous deployment of nodes, interference becomes a critical issue. Hence, the implementation of interference cancellation and suppression techniques will be essential for dense network applications.

Several interference suppression techniques make use of multiple receive antennas, combining their signals to isolate the desired signal from interference and noise. Hence, this is closely linked to the MIMO implementation of the simulator. Many interference suppressing receiver designs exist, some focusing on the suppression of interference between different spatial layers (Zero-Forcing (ZF), Minimum Mean Square Error (MMSE)), others on reducing inter-cell interference from neighbouring cells (Maximum Ratio Combining (MRC), Interference Rejection Combining (IRC)), and others combine both (Successive Interference Cancellation (SIC)).

While only ZF reception is implemented in the Vienna SL, the SINR computation is well suited to the creation of additional interference suppressing receiver types, as the signal, interference and noise components of the SINR can all be exposed individually and the receive filter response can be applied to each to obtain the post-equalisation SINR. We will see an example of this in Section 3.2.7.

In ns-3 interference suppressing filters are modelled implicitly as a set of predefined modifications to the SINR distributions. While this implementation has low computational complexity it is limited in flexibility as only a small number of predefined inter-layer interference cancelling configurations are implemented.

The simplified L2S interface of HetDenSim does not currently support this.

Radio Resource Management

Radio Resource Management (RRM) comprises strategies to assign and share radio resources in the network. These include inter-cell interference coordination techniques between neighbouring cells and scheduling time and frequency resources between users within the same cell.

Reuse and Inter-cell Interference Coordination One of the first strategies adopted to avoid interference between adjacent cells was to limit the reuse of spectrum; less interference is caused using higher spectrum reuse factors. However, this implies lower spectral efficiency in the network. Hence, the choice of spectrum reuse factor is subject to a trade-off between mitigating interference and spectral efficiency.

Re-examining spectrum reuse factors may be necessary for EDNs, as interference among neighbouring cells tends to increase with the density of base stations. For instance, LTE has been designed to work with low reuse factors (e.g., 1) in macrocell scenarios; however, it is likely that the reuse factor for EDNs will increase as a function of the cell density. Moreover, given that cells in EDNs are likely to have an irregular spatial deployment, it is not possible to use standard spectrum reuse patterns developed for hexagonal macrocell grids.

Simulators should allow variation of the spectrum assignment and reuse strategies depending on the cell density and distribution patterns. The Vienna SL simulator implements static frequency reuse and Frequency Fractional Reuse (FFR) for macrocell scenarios distributed in a hexagonal grid. HetDenSim implements these schemes, and in addition Frequency Aloha [81], which is the most basic form of spectrum reuse for small cells. Spectrum reuse through FFR has very recently been included in the most recent release of ns-3 also during the Google Summer of Code [82].

There are several other methods in which interference coordination can be carried out. For example, power control strategies and Coordinated Multi-Point (CoMP) can be considered as interference coordination techniques. While not yet present in ns-3 or HetDenSim, support for basic CoMP coordinated scheduling has been included in the most recent release of the Vienna SL simulator.

Scheduling Scheduling algorithms play an essential role in optimising different aspects of network performance, such as cell throughput or Quality-of-Service of all UEs. The choice of simulator to use for scheduler testing will depend on both the metrics to optimize and the types of feedback available for scheduler decisions, e.g. Channel State Information (CSI).

The considered SL simulators support schedulers for full-buffer traffic. HetDenSim has Round Robin (RR), Proportional Fair (PF), and Best CQI (B-CQI) scheduling, while Vienna SL additionally offers Max-Min Throughput. ns-3 provides these and more sophisticated schedulers for non-full buffer traffic, such as Channel and Quality-of-Service aware scheduling and Priority Set scheduling, which handles traffic for users with and without Guaranteed Bit Rate.

Traffic types

Both the volume of traffic and the number of traffic service types is expected to grow hugely in Extremely Dense Networks where heterogeneous devices and technologies coexist [76]. Simulators should accurately model various traffic types, while also providing the ability to create new traffic services.

If the volume of users does not increase at the same rate as the cell density, the average number of users per base station will tend to decrease. We can then infer that the smaller the size of the cell, the more dynamic the traffic will be, meaning that smaller cells might switch between active (i.e., users to serve) and inactive states frequently. This will cause time variations in the interference generated by the cell. For this reason, SINR distributions, which are usually obtained assuming full buffer traffic, will not give a full picture of network performance and should be complemented with other metrics, such as user perceived rate [83].

Of the simulators considered, the HetDenSim simulator currently only implements full buffer traffic, while the Vienna SL simulator, despite supporting multiple traffic type transmissions, models all inter-cell interference as full-buffer traffic. ns-3 includes several additional traffic types, such as voice traffic and multimedia; it is also possible to create new applications that generate traffic according to statistical models.

Upper layers

Packets usually go through many protocol layers before reaching their final destination. These layers, due to flow control and/or Automatic Repeat ReQuest (ARQ) mechanisms, will influence the delay and the amount of traffic transmitted through the air interface, which affects the user-perceived performance. For this reason, modeling these mechanisms is important when simulating EDNs. Moreover, to simulate some aspects of the network, such as handover performance, connection establishment times, performance of backhaul traffic, etc., functionality above the Medium Access Control (MAC) layer is required.

While system level simulators usually model only up to the MAC layer, network level simulators allow the user to investigate the system all the way up to the application layer. Among the SL simulators analysed, Vienna and HetDenSim do not model the higher layers of LTE. In the network level simulator considered, LTE

for ns-3 can be configured in two ways: with the Evolved Packet Core (EPC) or without it. With the EPC, higher layers such as Radio Link Control (RLC), Packet Data Convergence Protocol (PDCP), Radio Resource Control (RRC) and other protocols related to the core are modeled. Without it, a simplified version of the RLC allows the simulator to function with only full buffer traffic, disregarding the higher layers for performance purposes.

Backhaul

In Extremely Dense Networks, multiple base stations are connected to the core network through many heterogeneous backhaul technologies such as point-to-point fiber links, Passive Optical Networks (PONs), Digital Subscriber Line (DSL), microwave relays, etc. Frequently, studies of mobile networks assume an ideal backhaul, in which the fixed links have no delay or bandwidth limitations. In practice, the capacity and latency of the backhaul are constrained and this can degrade the performance of the wireless network.

It is possible to use the resources more efficiently, if the backhaul is shared among many base stations and able to adapt the capacity assured to them depending on the load.

To perform optical-wireless integration studies, where the MAC of the wireless and optical networks are designed together, EDN simulators must be able to model the backhaul. Due to the lack of higher layer protocols in system level simulators, network level simulators are better suited for this task. Among the evaluated simulators, ns-3 is the only simulator that can simulate the EPC and additional transport technologies such as simple point-to-point links or PONs [84].

Emulation

Even though simulation can give us valuable insights into the network performance, these do not replace practical prototypes with actual testbeds. Unfortunately, designing a full system is a complex task, not feasible in many research centers. For this reason, an important feature a simulator might have is the capability to connect to real hardware, while emulating other parts of the network [85]. Out of the simulators examined, only ns-3 is designed considering this possibility.

3.1.3 Comparative Analysis

In this subsection, we compare the suitability of the investigated simulators to EDN simulations from ease-of-use and extensibility perspectives. Furthermore, we provide a numerical comparison of how scalable the simulators are to EDN scenarios.

Ease-of-Use and Extensibility to EDNs

For the purposes of this work we define ease-of-use as the ease with which a user can set up and run a (dense network) simulation, making use of already supported functionalities and only altering parameter configurations. By extensibility, we refer to the ease of developing additional features for simulating previously unsupported scenarios.

To arrive at an objective assessment, we consider some simple ease-of-use and extensibility metrics and indicators, which we will use to compare the representative simulators.

Table 3.1: Ease-of-use Metric/Indicators

Simulator	Vienna SL	ns-3
Coding Language	MATLAB	C++ (ns-3)
No. of primary publication citations ¹	192	28
No. of forum posts within a month ²	35	64
Licensing	Free of charge for academic use	Open Source (GLP)
Predefined code testing script	Limited	Yes

Prerequisite Knowledge The knowledge required to use a simulator greatly affects its ease-of-use and can be a reason to prefer one simulator over another, either due

¹From scholar.google.com 25/04/2014.

²For the month 25/03/2014 to 25/04/2014. For Vienna SL all forum posts of the thread "System Level Questions" between these dates are included. As ns-3 is broader than just LTE only posts containing the keywords "LTE" or "LENA" are included.

to the detail with which particular parts of the system are modelled, or due to the level of programming knowledge required.

Regarding the programming skills required by different simulators, these are more a function of the programming language used and the design patterns applied than of the nature of the simulator. Among the simulators we consider for our investigation, Vienna SL and HetDenSim are MATLAB-based, while ns-3 is designed using C++. This makes Vienna SL and HetDenSim easier to learn, at the cost of processing speed. C++ usage is often regarded as more complex as it exhibits a number of complications which MATLAB handles transparently, e.g. memory management. Nonetheless, ns-3 makes use of many advanced design patterns, such as object factories, smart pointers and functors [86]; even though a user might not be familiar with all these concepts, knowing them can greatly simplify simulator usage and understanding.

Documentation Without proper documentation, learning how to use a new tool can be a daunting task. In the Vienna SL simulator, documentation is provided in the form of explanatory publications (e.g. [8]) containing system block diagrams, as well as listings of exposed parameters. In ns-3, documentation of the LTE model is partitioned into a model guide and user documentation. The model guide explains the details and reasoning behind each simulator feature, while the user documentation is aimed at aiding the user to set up, configure and examine their simulations. Additionally, documentation is automatically generated using Doxygen from source code comments. This tool enables users to check the purpose of classes and functions and simplifies code navigation.

User Community and Active Forum A helpful user community with an active forum can accelerate simulator familiarisation and clarify documentation ambiguities. The size of the community greatly influences the level of feedback. A large community increases confidence in underlying simulator functionality as problems are more likely to be discovered.

To quantify the community size and activity, Table 3.1 outlines some key metrics for the Vienna SL simulator [8] and the LTE module for ns-3 [78]; HetDenSim is omitted given its custom nature. Interestingly, while the main Vienna SL publication is more widely cited, the ns-3 LTE module receives almost double the Vienna SL's forum activity.

Open access to simulator source code allows the user community to correct errors and extend functionality. The ability of community members to independently distribute these extensions (open source) can greatly accelerate the rate at which new functionality develops and enables faster correction of inconsistencies.

Network Deployment Scenarios Simulation setup can be accelerated or hindered significantly depending on whether simulators support desired scenarios already and on whether they facilitate easy development of new scenarios. Several differences exist between the chosen simulators regarding this.

The Vienna SL simulator supports a wide range of predefined scenarios, mostly in keeping with the scenarios defined within the 3rd Generation Partnership Program (3GPP) standards. It is possible to define other scenarios, although use of non-default parameter combinations can often result in incompatibilities.

The NL simulator ns-3 offers a higher level of flexibility. Some predefined scenarios exist, but it is also easy to create new scenarios and develop or replace elements of scenarios.

The custom simulator has been designed with specific scenarios in mind, and though it is possible to simulate other scenarios, it might be necessary to create the components required.

Structure and Modularity When designing a complex system, a key concept that the designer must bear in mind is its structure and modularity. By isolating components according to their functionality, basic components can be made re-usable and their maintenance made easier. Object Oriented Programming (OOP) enables this modular design by encapsulating data and functionality into abstract data types.

As EDNs become more and more complex, the concept of modularity becomes increasingly important. In these scenarios, a single node may have multiple radio access technology interfaces, making the ability of a node to switch transmissions between one technology and another an important one. In the simulators evaluated, Vienna SL is focused on LTE, making it difficult to support other technologies. ns-3

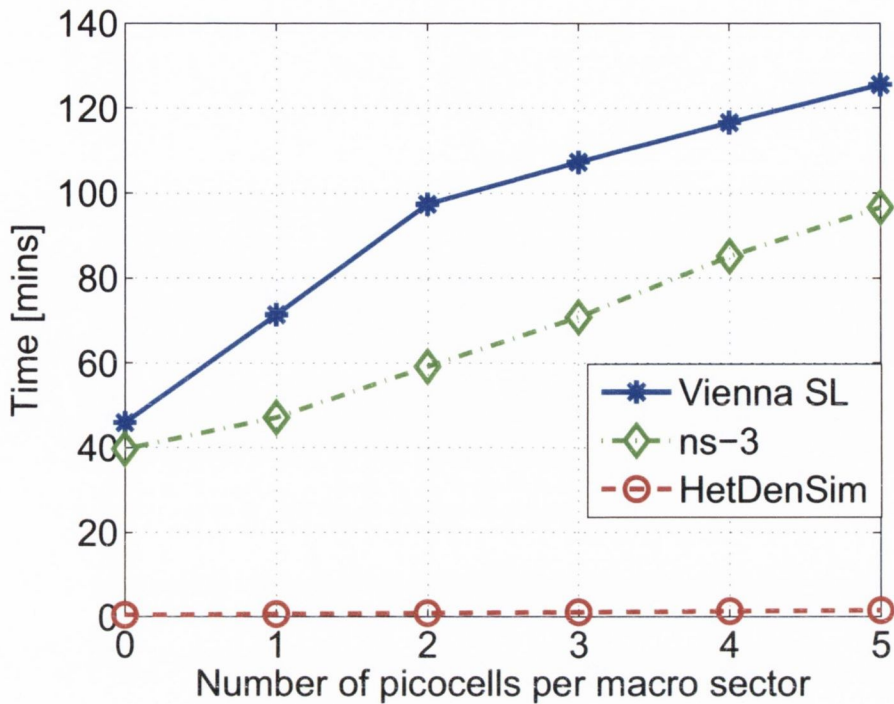


Figure 3.4: Variation in simulation time with small cell deployment density.

on the other hand, was developed as a general purpose simulator and thus naturally supports other technologies, such as 802.11.

Scalability for EDNs

To assess the suitability of the selected simulators for large scale network simulation, this section tests how simulation times scale with network size.

We investigate the 3GPP system level scenario for simulation of heterogeneous deployments [87] and vary the density of small cell deployment. The main parameters are reported in Table 3.2.

In Fig. 3.4, we show the run times of each simulator as we increase the density of low power nodes (i.e. picocells) per macrocell sector. The simulation time of the custom simulator, HetDenSim, is shown to be between 65 and 100 times less than that of ns-3, and between 85 and 115 less than that of Vienna SL. Moreover, HetDenSim's simulation times scale less steeply with network size than those of ns-3 or Vienna SL. Showing that through simplification of the simulator design, which may be acceptable

Table 3.2: Simulation Parameters and Hardware Configuration

Parameters	Value
Macro eNB deployment	57 macro eNB hexagonal grid, 3GPP case 1 [87, Table A.2.1.1-1]
Pico eNB deployment	Outdoor RRH/Hotzone, 3GPP case 6.2 [87, Table A.2.1.1.2-2], {0, 1, 2, 3, 4, 5} pico eNBs per macro eNB
UE deployment	3420 UEs uniformly distributed
Propagation model	3GPP Model 1 [87, Table A.2.1.1.2-3]
Fast fading ³	Rayleigh Fading (Pedestrian B)
UE Speed	3 Km/h
Macro eNB antenna pattern/gain	Directional antenna, 3GPP model [87, Table A.2.1.1-2] / 14 dBi
Pico eNB antenna pattern/ gain	Omnidirectional / 5 dBi
UE antenna pattern/ gain	Isotropic / 0 dBi
Bandwidth	10 MHz
Carrier Frequency	2.1 GHz
Frequency reuse	Full reuse 1
Traffic	Full buffer
Scheduler	Round Robin
Simulation length	250 TTIs
Hardware	Value
Processor	GenuineIntel i7-4930K CPU@3.4GHz [Family 6, Model 62, Stepping 4]
RAM	16 GB

for certain investigations, huge savings in run-time can be achieved

The lengthy simulation times of Vienna SL and ns-3 stem mostly from the detailed L2S interface, which despite slowing down the simulations, makes it possible to investigate effects of many MIMO techniques and fast-fading effects which are not modelled in HetDenSim. A few points should be noted; firstly, while Vienna SL and HetDenSim only simulate downlink, ns-3 also incorporates the uplink, thus increasing its run times. Secondly, for a fairer comparison, the EPC was not enabled in ns-3. The programming language in which the simulators were written, will also significantly affect the simulation duration. The more light-weight

³Fast fading is not performed by HetDenSim.

C++ (in which ns-3 is written) is generally faster than MATLAB (used for Vienna SL and HetDenSim).

This example demonstrates the tradeoff between simulation speed, ease-of-use, and the range, precision and types of supported functionalities, when choosing a simulator. If we target experiments on interference mitigation techniques, requiring an L2S interface with a low level of abstraction, then the choice should aim toward an SL simulator, such as the Vienna SL. Alternatively, if we need to assess RRM algorithms for large networks, then simulators with a less intensive PHY implementation like HetDenSim would be beneficial in terms of shorter simulation time. If our aim is to investigate network protocols, a NL simulator such as ns-3 is most suitable.

3.2 Extensions to the Vienna SL Simulator

As our work requires support for comprehensive simulation of multi-antenna transmissions at the system level we use the Vienna SL simulator (version 1.7r1119) as our primary baseline. In comparison, the simpler MIMO implementations provided by ns-3 and HetDenSim were not sufficient for our desired investigations, while the upper layer protocols only provided by ns-3 were not necessary for our investigations due to our assumptions of full buffer traffic and ideal backhaul information exchange between cooperating small cell eNBs. These assumptions were made to avoid overcomplicated simulation results which can potentially lead to misinterpretation of observed phenomena.

Despite the Vienna SL simulator providing a large range of already implemented functionalities, extensions to the functionality set were required for our desired investigations. These required extensions mainly centred around MU-MIMO implementation, including: MU-MIMO SINR and throughput computations, MU-MIMO scheduling, ZF MU-MIMO precoding, and MU-MIMO CSI feedback, as well as non-MU-MIMO-related extensions such as reassignments of UEs between neighbouring eNBs and computations of neighbouring cell CSI¹.

In this section we detail the extensions made, how they were implemented, and how we verified their performance.

¹It should be noted that since our implementation of these functionalities, a newer version of

3.2.1 Multi-User MIMO System Level Implementation

Our implementation of MU-MIMO is largely based on [17], in which the authors implemented LTE Rel. 10 ZF MU-MIMO at System Level in order to assess how it performs in a number of International Mobile Telecommunications-Advanced (IMT-A) standard scenarios. This implementation takes into account a number of very practical considerations, giving particular emphasis to the use of LTE standard-applicable feedback, such as: the use of Precoding Matrix Indicator (PMI) feedback to assess UE channel semi-orthogonality, and the use of the Maximum Expected SINR Combiner (MESC) method [88] to estimate CQIs for MU-MIMO scheduling and MCS selection.

In [17], emphasis is also placed on correct specification of channel overhead [89] taking into account control channel and reference signalling overhead, which is dependent on the number of Transmit (Tx) antennas and the maximum number of co-scheduled UEs for MU-MIMO.

Our MU-MIMO implementation was validated against the Cell Spectral Efficiency (CSE) and Cell Edge User Spectral Efficiency (CEUSE) results provided within the paper in [17, Table II], making use of the same full set of simulation configuration parameters as they did, although within the Vienna SL simulation environment.

Within the domain of the Vienna SL simulator, for this to operate correctly a number of changes to the Vienna core simulator operation were required:

Firstly, in order to reduce the size of the channel trace required within the simulation, a single long fast fading channel trace was formerly used, from which sub-traces were extracted randomly for each simulated UE. For Single-User MIMO (SU-MIMO) scenarios this works fine and reduces the computational burden on the simulator substantially; however, in MU-MIMO scenarios this results in unrealistic correlations between the channels of different UEs, reducing their likelihood of orthogonality. In order to remedy this, independent fast fading channel traces needed to be generated for each UE-to-eNB instead. This makes the

the Vienna SL simulator (v1.8r1375) has been released. For this Section 3.1 has been appropriately updated. In v1.8r1375, basic MU-MIMO and CoMP coordinated scheduling functionality were included; however, these represent simplified initial implementations and do not provide the level of detail required for our investigations. Instead, we plan to combine our extensions with theirs in the near future and to make our new-version-compatible extensions available for future similar investigations.

Table 3.3: SINRs used in MU-MIMO Simulation (single cell implementation)

UE-SINR	BLER-SINR
$\frac{ \mathbf{g}_k \mathbf{H}_{k,O} \mathbf{w}_k ^2}{ \mathbf{g}_k \sum_{l=1}^{N_e N_B - 1} a_l \mathbf{H}_{k,l} \mathbf{W}_l ^2 + \sigma^2 \mathbf{I} \ \mathbf{g}_k\ ^2}$	$\frac{ \mathbf{g}_k \mathbf{H}_{k,O} \mathbf{w}_k ^2}{ \mathbf{g}_k \sum_{j=1}^{n_{MU} - 1} \mathbf{H}_{k,O} \mathbf{w}_j ^2 + \mathbf{g}_k \sum_{l=1}^{N_e N_B - 1} a_l \mathbf{H}_{k,l} \mathbf{W}_l ^2 + \sigma^2 \mathbf{I} \ \mathbf{g}_k\ ^2}$
Computed by UE	Computed by simulator
Computed for UE feedback <ul style="list-style-type: none"> • Signal power based on best single-user precoder. • Does not include MUI. • Neighbour cell precoding based on single-layer random beam. Will later be adjusted to take into account MUI.	Computed post-scheduling, taking into account scheduling outcomes, including: <ul style="list-style-type: none"> • Final selected UE signal precoding, • Precoding of co-scheduled UEs, • Neighbour cell selected precoding.
For each RB	Averaged over scheduled RBs

simulation slower, but ensures that undesired inter-UE channel correlations are removed.

Secondly, an alteration to the link-to-system mapping was required which consisted of separating the SINRs used for UE feedback computation (UE-SINR) and the simulator SINR used in BLER and throughput calculations (BLER-SINR). Previously the BLER-SINR was obtained by performing an Mutual Information Effective SINR Mapping (MIESM)-average of the subband UE-SINRs over the Resource Blocks (RBs) in which the UE is scheduled. However, in MU-MIMO scenarios the UE cannot know what Multi-User Interference (MUI) it will experience prior to the scheduling process and so, despite this being required in the BLER-SINR, it cannot be included in the UE-SINR. Hence the BLER-SINR was separated out and altered to take into account, not only which RBs each UE is scheduled in, but also the levels of MUI resulting from the selected precoders of co-scheduled UEs. Subsequently, the most recent release of the Vienna SL simulator (version 1.8r1375) has performed a similar alteration. This is summarised in Table 3.3.

Additionally, a number of minor changes such as in the way scheduler feedback was specified were also required.

3.2.2 MU-MIMO User Selection Algorithm

In this work, when MU-MIMO is used, we schedule UEs according to the well-known MU-MIMO Semi-orthogonal User Selection algorithm [27], with some minor adaptations. This algorithm provides a heuristic solution to the MU-MIMO user selection problem which removes the need to do an exhaustive search of all possible UE pairing combinations. Also, the semi-orthogonality constraint that the algorithm employs ensures that the interference between co-scheduled UEs is kept below a certain level. Further, this algorithm is designed to operate with limited feedback channel direction information which is necessary when considering realistic implementations such as LTE scenarios.

The Semi-orthogonal User Selection (SUS) algorithm starts off by selecting the UE with the highest SINR, to be scheduled. Next, a set is computed of the unscheduled UEs which have quantised channel vectors semi-orthogonal to that of the already scheduled UE. From this UE set, the one with the highest SINR is selected to be co-scheduled for MU-MIMO transmission with the previously scheduled UE. This process is performed iteratively until either no more semi-orthogonal UEs remain or the maximum number of UEs that can be co-scheduled at once is reached.

Elaborating on earlier, two quantised channels vectors are called semi-orthogonal if the magnitude of the spatial correlation between them is below a predefined threshold, ϵ . The quantised channel vector for the k^{th} UE, $\hat{\mathbf{h}}_k$, is obtained from the feedback PMI as the pseudo-inverse of the recommended precoding vector. This vector forms a quantised estimate of the directionality of the channel. For this reason the PMI can also be referred to as the Channel Direction Indicator (CDI).

In our implementation we deviate slightly from [27] in the following ways: firstly, instead of selecting UEs based on their SINR we select them based on their proportional fair metric; secondly, after selecting UE pairs we perform an additional check to ensure that the proportional fair metric achieved through use of MU-MIMO is higher than the proportional fair metric of Single-User MIMO beamforming on a single spatial stream (SU-MIMO) of the initially scheduled UE on its own; if it is not, the initial UE is scheduled for SU-MIMO alone. This is of particular importance when we consider UEs which were reassigned from adjacent cells.

The proportional fair metric for the k^{th} UE, when co-scheduled as one of n UEs in RB(/subband) b , is computed as its expected instantaneous throughput $r_{n,k}^b$ divided by its long term average throughput R_k . R_k is obtained as an exponential average of past throughputs as

$$R_k(t+1) = \alpha R_k(t) + (1-\alpha)r_{n,k}^b \quad (3.2)$$

where $0 \leq \alpha \leq 1$ determines the rate at which the exponential average decays.

The modified user selection process for RB b is summarised in Algorithm 1 (PFSUS(b)) for our considered case of $N = 2$ without loss of generality. \mathcal{A}_n denotes the set of UEs with quantised channels semi-orthogonal to the currently scheduled UEs. Initially, this set contains all K UEs in the cell, as no UEs are yet scheduled.

Algorithm 1 PFSUS(b)

Initialise $\mathcal{A}_0 = \{1, \dots, K\}$

$$\pi_b(1) = \operatorname{argmax}_{k \in \mathcal{A}_0} \frac{r_{1,k}^b}{R_k}$$

$$\mathcal{A}_1 = \{1 \leq k \leq K : |\hat{\mathbf{h}}_k, \hat{\mathbf{h}}_{\pi_b(1)}| < \epsilon\}$$

$$\pi_b(2) = \operatorname{argmax}_{k \in \mathcal{A}_1} \frac{r_{2,k}^b}{R_k}$$

if $\sum_{j=1}^2 \frac{r_{2,\pi_b(j)}^b}{R_{\pi_b(j)}} > \frac{r_{1,\pi_b(1)}^b}{R_{\pi_b(1)}}$ **then**

Schedule $\pi_b(1)$ and $\pi_b(2)$ for MU-MIMO transmission

else

Schedule $\pi_b(1)$ for SU-MIMO transmission alone

end if

The scheduling algorithm is performed independently by each eNB and is repeated for each RB, beginning with the RB containing the UE of highest proportional fair metric and then iterating over all others. The average throughput is recomputed at each iteration. This ensures fast convergence to a proportional fair equilibrium [90] which is necessary when considering short simulation snapshots and is shown in Algorithm 2.

Algorithm 2 SUS iteration over RBs

Initialise $\mathcal{B} = \{1, \dots, B\}, \mathcal{A}_0 = \{1, \dots, K\}$
while $\mathcal{B} \neq \emptyset$ **do**
 $b = \operatorname{argmax}_{b \in \mathcal{B}, k \in \mathcal{A}_0} \frac{r_{1,k}^b}{R_k}$
 PFSUS(b)
 Recompute Eqn. (3.2)
 $\mathcal{B} = \mathcal{B} \setminus b$
end while

3.2.3 Zero-Forcing MU-MIMO

In cases where the quantised channel vectors of co-scheduled UEs are not fully orthogonal (for example if $\epsilon > 0$ is used in the above scheduling algorithm) Zero-Forcing Beam Forming (ZFBF) [22] can be applied to minimise interference leakage between the co-scheduled transmissions. This method computes unit-norm precoding vectors, $\mathbf{w}_{\pi_b(i)}$, for each co-scheduled UE $\pi_b(i)$, $i \in \{1, \dots, n_{MU}\}$, such that the computed precoding vector is orthogonal to the fed back quantised channel vectors of all other co-scheduled UEs

$$\hat{\mathbf{h}}_{\pi_b(j)} \mathbf{w}_{\pi_b(i)} = 0, \quad \forall j \neq i, \quad j \in \{1, \dots, n_{MU}\}. \quad (3.3)$$

This is performed independently for each RB $b \in \mathcal{B}$.

This operation proceeds as follows: a matrix $\hat{\mathbf{H}}$ is constructed by stacking the quantised channels vectors of co-scheduled UEs as

$$\hat{\mathbf{H}} = [\hat{\mathbf{h}}_{\pi_b(1)}^T, \dots, \hat{\mathbf{h}}_{\pi_b(n_{MU})}^T]^T. \quad (3.4)$$

Of this, the pseudo-inverse is taken to obtain the transmit precoding matrix

$$\mathbf{W} = \hat{\mathbf{H}}^\dagger = \hat{\mathbf{H}}^H (\hat{\mathbf{H}} \hat{\mathbf{H}}^H)^{-1}. \quad (3.5)$$

The precoding vector of UE $\pi_b(i)$, $\mathbf{w}_{\pi_b(i)}$, is then obtained as the normalised form of the i^{th} column of \mathbf{W} . Also, while the nulling of interference in ZFBF can result in reduced signal power in the direction at the desired UEs, in cases where the selected UEs are semi-orthogonal, this reduction is small.

As pointed out in [17], the use of ZFBF in LTE-Advanced (LTE-A) Rel. 10 systems is made possible by the Demodulation Reference Signal (DM-RS). This reference signal is precoded in the same way as the data transmissions and as such enables the receiver to compute from it the precoding vector that was used, which is necessary for successful demodulation.

It should also be restated, however, that due to the quantisation errors in the feedback PMIs used to construct the ZF beamforming vectors, this solution is not perfect and is not capable of fully removing MUI from transmissions.

3.2.4 MU-MIMO CQI calculation

In LTE systems, MU-MIMO usage is performed transparently to the UE, meaning it is unaware of whether it will be scheduled for SU- or MU- MIMO in subsequent transmissions. As a result UEs provide the same (SU-MIMO) PMI and CQI feedback regardless of whether SU- or MU- MIMO transmissions are to be performed [20, 91]. In LTE Rel.8 systems this results in reduced feedback overhead and computing complexity during MU-MIMO operation, while in LTE-A Rel.10 systems this is an even more important feature as it allows for seamless switching between SU- and MU- MIMO, which can also be on a subband basis.

As a consequence, for user scheduling and MCS selection purposes, it is necessary for the eNB *to convert* the SU-MIMO CQI feedback provided by the UE to an estimated MU-MIMO CQI. In the literature many methods of CQI estimation for MU-MIMO scenarios exist; however, often these work on the basis that the UE would feed back MU-MIMO-specific CQI instead of SU-MIMO CQI [27, 88] and so are not directly applicable for LTE systems. Further, the MU-MIMO CQI specification techniques in [27, 88] assume that the maximum possible number of UEs that can be co-scheduled, N_{MU} , will always be scheduled. In our case, however, we consider scenarios in which the number of UEs per cell is low, and as a result this assumption cannot be applied (if the channel conditions do not suit or there are too few UEs in the cell, there may instead be only a single UE scheduled for SU-MIMO). For this reason, and due to its compliance with the usage of SU-MIMO CQI feedback, we apply the MU-MIMO CQI adjustment method proposed in [20].

In this method, all UEs feed back CQIs for SU-MIMO (not including MUI)

which are adjusted to account for the splitting of the transmit power between the co-scheduled UEs, and for the mean level of unsuppressed MUI, Δ_{MUI} , as follows:

$$CQI_{MU-MIMO} = \frac{1}{\frac{n_{MU}}{CQI_{SU-MIMO}} + \Delta_{MUI}(n_{MU} - 1)} \quad (3.6)$$

where n_{MU} is the number of UEs co-scheduled for MU-MIMO within a single time-frequency RB, and $CQI_{SU-MIMO}$ is the fed back quantised SINR, in linear form.

Δ_{MUI} is a precomputed parameter, specific to the environment and device configuration. It is obtained by taking the expectation over a large number of channel realisations of the ratio of the suppressed MUI power to the signal power for SU-MIMO (without power splitting between co-scheduled UEs):

$$\Delta_{MUI} = \mathbb{E} \left[\frac{|\mathbf{g}_k \sum_{j=1}^{n_{MU}-1} \mathbf{H}_{k,O} \mathbf{w}_j|^2}{n_{MU} |\mathbf{g}_k \mathbf{H}_{k,O} \mathbf{w}_k|^2} \right]. \quad (3.7)$$

A fixed value of Δ_{MUI} is made reasonable by the constraint that the maximum allowed channel correlation of simultaneously served UEs in this thesis, ϵ , is very low. This means that between different effective channel matrix realisation pairs the variance in the amount of suppression that can be achieved is also low.

As Δ_{MUI} is computed using signal and interference channel samples generated as part of repeated system level simulations, we validated our Δ_{MUI} computations (obtained using the Vienna SL simulator) through comparison with the Δ_{MUI} Cumulative Distribution Function (CDF) provided in [20, Fig. 1]. As the antenna and receive filtering configurations used later in this thesis are similar to those used in [20], it is not surprising that they too share the same mean value of $\Delta_{MUI} = 0.05$.

3.2.5 Neighbouring eNB Feedback

In order to assess the feasibility of potential reassignments, when the reassignment mechanism is initiated a UE needs to check the quality of the channels between it and any target neighbouring eNBs. To this end, as detailed later in Sections 4.3 and 5.3, unpaired UEs compute a wideband CQI for the target cell and a wideband PMI.

As was discussed in Section 2.1.3, UEs are capable of computing channel estimates for neighbouring eNBs using Cell Specific Reference Signals (CS-RSs). CQI values

obtained in this way are then adapted for different values of n_{MU} in the same way as outlined in the previous subsection, using Eqn.(3.6).

For CoMP applications, in the 3GPP document [37], feedback of this sort is identified as Implicit Channel State Feedback, subject to the hypothesis of MU-MIMO usage on n_{MU} layers. As we will consider cases in which, after reassignment, the original cell may be switched to an idle state if there are no remaining UEs to serve, this will be taken into account through the additional hypothesis of original cell deactivation in cases where this is certain to occur (e.g. Ch5), in other words, with interference from the original cell removed.

This CSI is used both in a reassignment decision made at the UE side initially and, depending on the outcome of the first decision, subsequently shared with the centralised coordinator in the second stage of the decision. This feedback requirement is similar to in CoMP Coordinated Scheduling/Beamforming techniques in which the wideband PMI and an expected CQI improvement are shared with the centralised unit [40]. In our implementation the sharing of this information is assumed to be ideal.

3.2.6 MU-MIMO Spectral Efficiency

In both the scheduling and reassignment processes it is necessary to compare the expected relative performances of SU-MIMO and MU-MIMO. For this, we map the $CQI_{SU-MIMO}$ and $CQI_{MU-MIMO}$ values to expected rates per RB based on a fitting to Bit-Interleaved Coded Modulation (BICM) curves, as discussed in [92]. Essentially this is similar to taking the Shannon capacity, but also taking into account modulation and coding. We represent the expected rate of a UE k , when co-scheduled as one of n_{MU} UEs, attached to eNB e as $r_{n_{MU},k,e}$. In this work e will be either the currently attached eNB, which we will call the original eNB, O , or a neighbouring eNB targeted for reassignment, T .

When MU-MIMO is used ($n_{MU} \geq 2$), a single RB is shared amongst multiple co-scheduled UEs, each on different spatial layers; as the UEs each only part-occupy the RB bandwidth, their MU-MIMO spectral efficiency in bits per RB is given by their rate for the RB, multiplied by the number of simultaneously served UEs that the RB is divided between.

3.2.7 MMSE-IRC Filtering

Minimum Mean Square Error (MMSE) with Interference Rejection Combining (IRC) filtering combines the suppression of co-layer intra-cell interference by the MMSE filter and suppression of inter-cell interference through IRC filtering. This receiver is also known as the Advanced LTE UE Receiver and has been defined as the new baseline receiver from 3GPP LTE Rel.11 onward [53]. This is of particular benefit in dense small cell scenarios such as those we investigate.

For system level simulation, the MMSE-IRC receive filter can be expressed as [93]

$$\mathbf{g}_k = \mathbf{h}_{\text{eff},k}^H (\mathbf{h}_{\text{eff},k} \mathbf{h}_{\text{eff},k}^H + \mathbf{C}_{\text{MUI}} + \mathbf{C}_{\text{INT}} + \mathbf{C}_{\text{W}})^{-1}. \quad (3.8)$$

Here, we use \mathbf{h}_{eff} to represent the effective channel vector, defined as the product of the channel matrix \mathbf{H} and the precoder \mathbf{w} (in other words $\mathbf{h}_{\text{eff},k} = \mathbf{H}_{k,O} \mathbf{w}_k$). \mathbf{C}_{MUI} , \mathbf{C}_{INT} and \mathbf{C}_{W} represent the multi-user, inter-cell and white noise interference covariance matrices, respectively:

$$\mathbf{C}_{\text{MUI}} = \sum_{j=1}^{n_{\text{MU}}-1} \mathbf{h}_{\text{eff mui},j} \mathbf{h}_{\text{eff mui},j}^H, \quad \mathbf{h}_{\text{eff mui},j} = \mathbf{H}_{k,O} \mathbf{w}_j, \quad (3.9)$$

$$\mathbf{C}_{\text{INT}} = \sum_{l=1}^{N_{\text{cNB}}-1} \mathbf{H}_{\text{eff int},l} \mathbf{H}_{\text{eff int},l}^H, \quad \mathbf{H}_{\text{eff int},l} = \mathbf{H}_{k,l} \mathbf{W}_l, \quad (3.10)$$

$$\mathbf{C}_{\text{W}} = \sigma^2 \mathbf{I}, \quad (3.11)$$

where $\mathbf{h}_{\text{eff mui},k}$ and $\mathbf{H}_{\text{eff int},l}$ are the effective channel vectors(/matrices) for the multi-user and inter-cell interference.

3.2.8 Reassignment

In the case where UE k is reassigned to a neighbouring cell, a number of elements in the SINR expression will change. Most significantly, the signal channel from the original cell $\mathbf{h}_{k,0}$ will become interference and the interference from the target cell $\mathbf{h}_{k,T}$ will become the signal channel, where T is the target neighbouring cell into which

UE k gets reassigned. As a result, the precoders and receive filters used \mathbf{w}_k , \mathbf{w}_j , \mathbf{W}_l and \mathbf{g}_k will all require recomputing to correspond to the new channel vectors.

We do not implement a reassignment delay, meaning that transmissions to the reassigned UE originate from the new cell in the Transmission Time Interval (TTI) following directly following the reassignment.

3.3 Indoor and Outdoor Scenarios

In order to enable reliable performance comparison between different wireless systems, organisations such as 3GPP and IMT-A make use of a number of standard deployment scenarios which attempt to cover all major deployment types. For example, the IMT-A-defined baseline scenarios include [94]: Urban Macrocell (UMa), Suburban Macrocell (SMa), Rural Macrocell (RMa), Urban Microcell (UMi) and Indoor Hotspot (InH). However, while for the majority of these scenarios a simple hexagonal grid provides a good fit (e.g. UMa, SMa, RMa, UMi), in small cell scenarios such a model cannot be applied. This results largely from two main issues: small cell scenarios tend not to follow a regular deployment pattern; and walls and other objects substantially affect signal propagation when the distances considered are small. Unfortunately, the inclusion of walls means that the produced results can become very scenario specific. Further, given that the concept of heterogeneous networks and the scenarios in which small cells are envisioned are constantly evolving, simulation scenarios for these types of deployment are constantly being updated. Hence, for small cell/HetNet scenarios, the range of possible baseline deployment cases is quite large and diverse. For example [94] defines the InH scenario, [95] defines Dual Stripe and Suburban HeNB models, [96] defines a set of scenarios labelled Virtual reality office, Madrid grid, shopping mall and stadium, and [87] defines five different scenarios labelled cases 5.1, 5.2, 5.3, 6.1 and 6.2. As a result the choice of which small cell scenario to use potentially becomes a very important decision.

With this in mind, we explore the potential benefits of the proposed coordinated MU-MIMO-based reassignment and sleep mode activation mechanisms in two small cell scenarios, one indoor and one outdoor, both subject to external macrocell interference. The first is the 3GPP Dual Stripe indoor residential femtocell




Figure 3.5: Dual Stripe scenario SINR distribution from simulation with a DR of 0.2.

deployment, while the second is an outdoor picocell deployment based on real base station locations in Stachus Square, Munich city centre. Both of these scenarios were implemented by us in the Vienna SL simulator as part of this work.

3.3.1 Indoor Residential Scenario

For investigating indoor residential femtocell (HeNB) deployments the Dual Stripe model [95] is the most widely used scenario. It provides a realistic, yet reproducible, simulation scenario and also provides a sufficient level of flexibility over both base station and user deployment densities. This model consists of two apartment buildings side-by-side, each subdivided into apartments separated by walls. In this work we consider single-story deployment in which each of the two apartment buildings contains two rows of 10 apartments, meaning a total of 40 apartments. The density of HeNBs deployed is indicated by the Deployment Ratio (DR), which, more specifically, denotes the probability of a given apartment containing a HeNB. So, for example, with 40 apartments and a DR of 0.2, on average of 8 apartments will contain a HeNB.

External interference is experienced from a tri-sector macrocell base station located 60m to the left of the bottom-left corner of the apartment building. The

three sectors of this macrocell interferer each transmit a power of 40W (46dBm) while the femtocell base stations each transmit at 65mW (20dBm).

From a coordination point of view, this scenario is similar to already-centrally deployed Ericsson Radio Dot systems [69]. In large buildings, through use of software defined radio, these deployments allow for coordination of the deployed small cells on each floor [97].

The SINR distribution of an example deployment is illustrated in Figure 3.5. White dots mark the HeNB locations, while the orange dot marks the location of the external macrocell.

This scenario will be used in both Chapters 4 and 5.

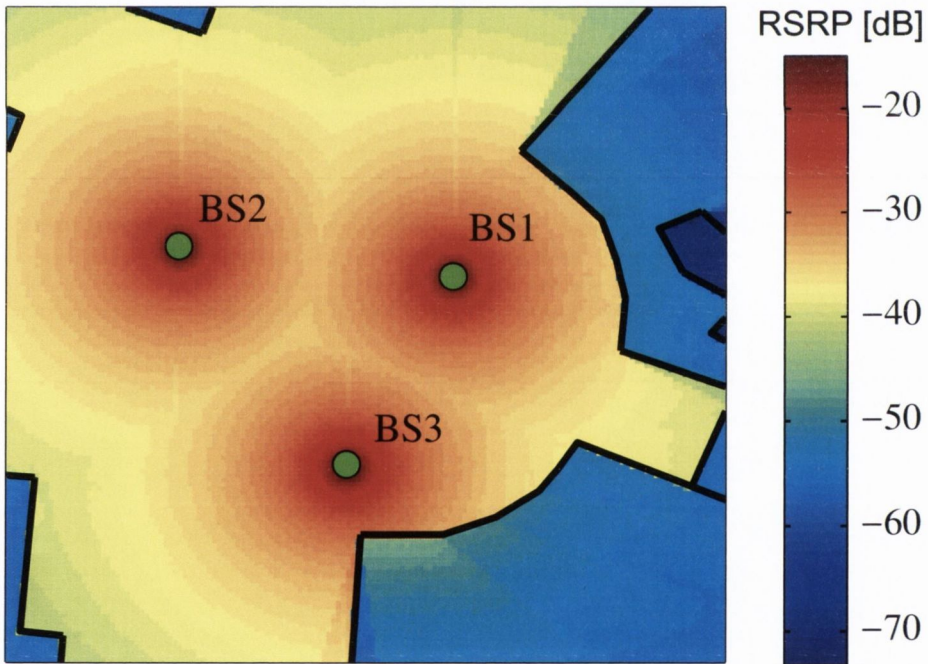
3.3.2 Outdoor City Centre Scenario

Within our research group there is a strong emphasis on the usage of real cellular network data due to its unrivalled ability to demonstrate the feasibility of proposed solutions. This has included usage of deployment and traffic data from two of Ireland's top network operators in assessing the effectiveness of network sharing solutions [98], usage of operator deployment data from a range of Polish cities for coverage provisioning analysis [99], as well as usage of an extensive publicly available set of RF spectrum measurements [100] in assessing how well reinforcement learning techniques for dynamic channel selection actually perform in real cases [101]. In this case, through collaboration with Bell Laboratories Alcatel-Lucent, Dublin, we have gained usage of a set of detailed receive power maps covering the Stachus Square area of Munich City centre.

The outdoor environment Munich city centre received power maps were obtained using the Wireless System Engineering (WiSE) 3D ray-tracing tool [102]. This tool takes into account building locations and compositions and applies reflections, refractions, diffusions and propagation to emitted base station transmissions in order to accurately estimate received signal powers in real-life environments. This provides us with a scenario which comprehensively and realistically takes into account pathloss, shadow fading and penetration losses. A detailed discussion of 3D ray-tracing and how it is used to model wireless propagation environments can be found in [38, Section 14.2.3].



(a)



(b)

Figure 3.6: (a) Stachus Square (b) Strongest received signal power from the three picocells.

Within this scenario we consider two different small cell deployment cases: one with 3 picocells and one with 21 picocells.

In both deployment cases the macrocell site locations, sector orientations and antenna gains correspond to those of one of the top-tier mobile operators in Germany, while the picocell locations were chosen to provide good coverage in the desired areas. All picocell base stations are equipped with omnidirectional dipole antennas and transmit at a power of 250mW (24dBm), while interfering macrocell base station sites each consist of 3 sectors, each transmitting at a power of 20W (43dBm).

3-Picocell Deployment Case

Figure 3.6a shows a photo of the square in Munich on which this scenario focuses. In this deployment case 3 picocell base stations (marked by green dots) serve users within a 150m x 150m area, while macrocell interference is observed from 4 surrounding tri-sector base station sites (blue dots). The strongest Reference Signal Received Powers (RSRPs) from the 3 small cell base stations, as computed by the WiSE 3D ray-tracing tool are shown in Figure 3.6b.

This scenario will be used in Chapter 4.

21-Picocell Deployment Case

Figure 3.7 shows the strongest picocell RSRP from the 21 picocell base stations in this case, at each point within the 250m x 250m investigated Region of Interest (ROI). The locations of the 21 picocell eNBs are marked by cyan dots. In this case, UEs experience macrocell interference from 6 surrounding tri-sector base stations with locations, again, corresponding to those of one of Germany's top-tier mobile operators.

This scenario will be used in Chapter 5.

3.4 Summary

In this chapter we have outlined a number of key simulation considerations which we undertook in order to simulate the small cell reassignment mechanisms proposed in

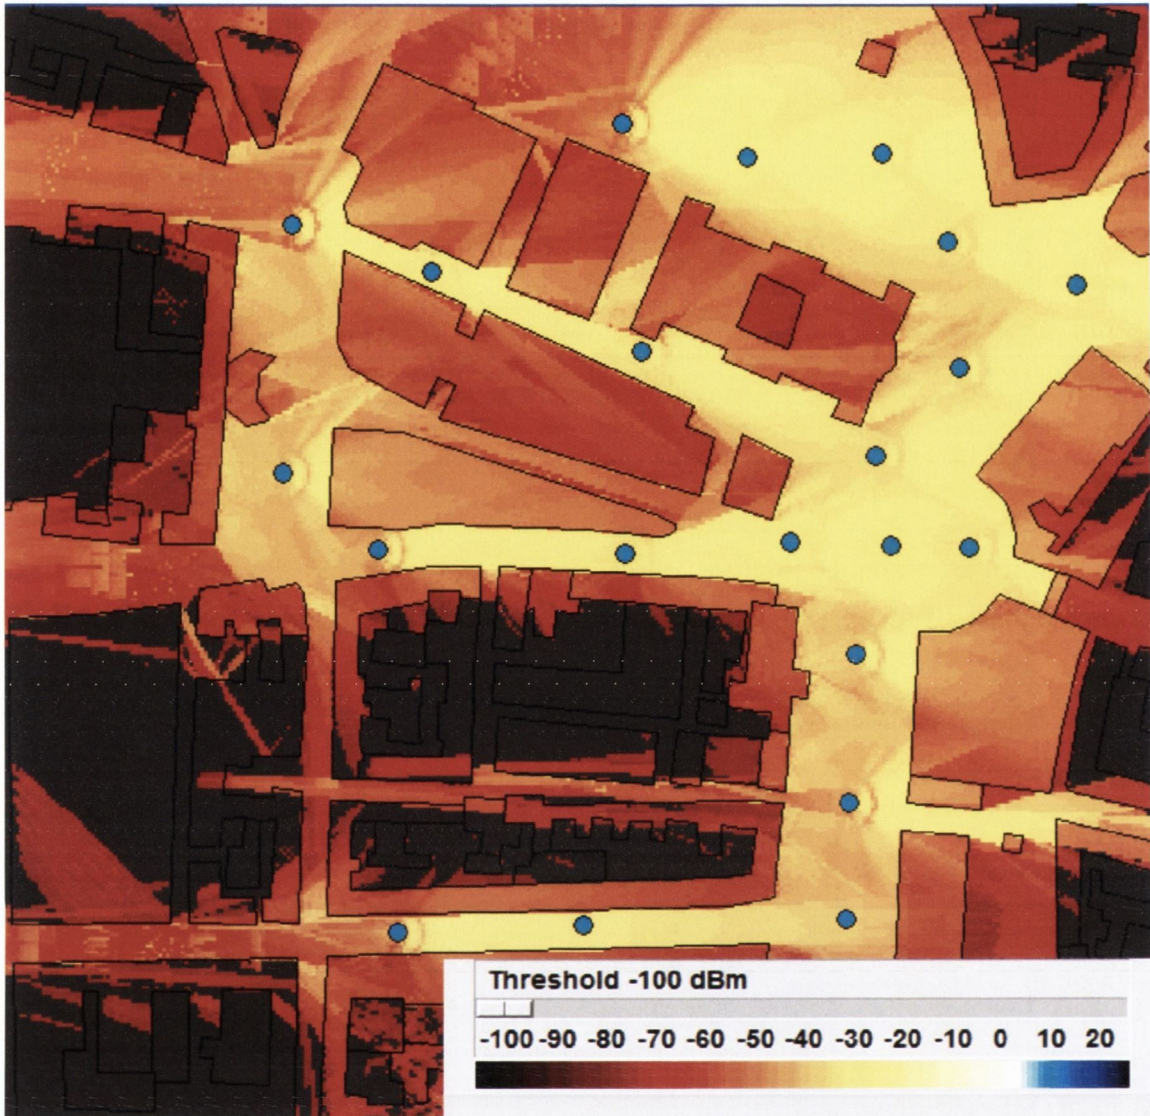


Figure 3.7: Maximum RSRP map for Stachus Square, Munich city centre, 21-pico cell deployment.

the following two chapters.

We compared a number of LTE simulators, suitable for small cell system/network level simulations, to use as a baseline. We compare these both qualitatively and quantitatively from perspectives such as supported features, ease-of-use, and scalability for dense networks. From these we selected the Vienna SL simulator mainly due to its ability to comprehensively support multi-antenna scenarios. With this as a baseline, we have detailed a number of key extensions to the Vienna SL simulator which we implemented as part of this work and explained

how the underlying functionalities of these extensions operate. Finally, we have discussed the indoor and outdoor deployment scenarios which are used in the following chapters and our motivations behind their selection.

In the following two chapters these developed simulation capabilities and environments will be used to analyse the performance our proposed MU-MIMO-based reassignment mechanisms in term of increases in spectral efficiency and MU-MIMO usage, as well as decreases in energy consumption of the network.

4 MU-MIMO Across Small Cells: Spectral Efficiency Increases

This chapter is largely based on the results and contributions of our paper entitled "*Improved Spectral Efficiency through Multi-User MIMO Across Small Cells*" [103] which was submitted to the *IEEE Transactions on Vehicular Technology*, and our conference publication, published in the *IEEE International Conference on Communications* entitled "*Multi-User MIMO Across Small Cells*" [104], which it extends.

4.1 Introduction

The main contribution of this chapter is centred on the assessment of our proposed *Multi-user MIMO across Small Cells* concept. In small cell networks, where the number of User Equipments (UEs) per cell is low, finding suitable sets of UEs to be co-scheduled for Multi-User MIMO (MU-MIMO) is not always possible. In these cases we propose MU-MIMO-based cell reassignments of users into adjacent cells to enable MU-MIMO operation. This chapter investigates the gains of this new concept for realistic indoor and outdoor small cell scenarios, taking into account user scheduling before and after the reassignment occurs. From system level simulations we found that the proposed mechanism results in considerable increases in spectral efficiency and MU-MIMO usage by both the reassigned UE and the UE in the new cell that it gets paired with. We also found that, in the outdoor small cell scenario, the percentage of UEs that meet the reassignment criteria and can benefit from neighbour-cell MU-MIMO is almost double that of the indoor scenario.

The level of emphasis on MU-MIMO operation in Long Term Evolution (LTE) systems increases with each new release of the standards, from basic support in the initial release, to combined Single-User MIMO beamforming on a single spatial stream (SU-MIMO)/MU-MIMO operation in release 10, toward 3D beamforming and massive Multiple Input Multiple Output (MIMO) in future releases [1], where UEs are co-scheduled on different downtilts and where the number of antennas at the base station is increased by orders of magnitude to enable enormous numbers of UEs to be coscheduled at once, respectively.

In parallel, advances in small cell technology are developing at a similarly fast rate [1] and their deployment in both indoor and outdoor scenarios is growing rapidly each year, having overtaken the total number of macrocells worldwide between October and November 2012 [105].

Nonetheless, the combination of MU-MIMO and small cell technology is not often investigated. This is because it is often simply accepted that, seeing as it is more difficult to consistently find good MU-MIMO groupings when there are fewer UEs in the cell, MU-MIMO should mainly be limited to cases where the number of UEs is high, i.e. macrocell scenarios. In fact, MU-MIMO works commonly consider 20 [88, 106], 50 [2] or even 100 [27] UEs per cell when assessing potential system gains. Two works which do investigate MU-MIMO usage in small cell environments can be found in [61] and [62]. In these it is found that the gains of MU-MIMO are severely limited when there are few UEs in each small cell and that as the deployment density increases MU-MIMO performance is limited even further. However, if good MU-MIMO pairings can be found when the number of UEs served is low then the full benefits of MU-MIMO can still be obtained. In order to achieve this we define a new network reconfiguration/coordinated multi-cell scheduling approach, *MU-MIMO Across Small Cells*.

In this work, we consider small cell scenarios in which the number of UEs in each cell is low, meaning that the likelihood of finding UEs in the same cell with orthogonal precoding matrices is also low. At the same time, however, small cell scenarios often exhibit strong line-of-sight components, leading to high correlation between transmit antennas, which is well suited to MU-MIMO operation. Further, due to the dense deployment of these networks, UEs often have sufficiently high Signal to Interference and Noise Ratios (SINRs) in neighbouring cells to support MU-MIMO operation. For this reason, we propose the reassignment of UEs to neighbouring cells in order

to create pairings where they were previously unavailable and obtain the benefits of MU-MIMO across multiple small cells.

In a recent survey on multi-cell scheduling in LTE/LTE-Advanced (LTE-A) systems [40], the authors subdivide the topic into the categories of: Inter-Cell Interference Coordination; Coordinated Multi-Point (CoMP); and the crossover between the two, namely Coordinated Beamforming/Scheduling. In the following few paragraphs we will outline the fundamental differences between our work and Coordinated Beamforming/Scheduling, as well as the jointly processed CoMP concepts of Network MIMO/Cloud Radio Access Network (C-RAN) and Dynamic Cell(/Point) Selection.

Of the three categories identified in [40], the one which most closely relates to our work is the crossover category of Coordinated Beamforming/Scheduling [42, 46]. Here, UEs get scheduled with the precoders that cause the least interference to the UEs operating on the same frequency in their neighbouring cells. In our work, however, the emphasis is different. Instead of scheduling UEs with orthogonal beam patterns in their respective cells, which has the effect of reducing interference, we reassign one of the UEs between cells so that MU-MIMO can be used, from which we can achieve an increase in capacity through use of additional spatial streams.

Another somewhat similar concept, which, as pointed out in [107], can be seen as a generalised form of MU-MIMO, is Network MIMO. Network MIMO is basically MU-MIMO but with the transmit antennas spread over multiple base stations (eNBs) rather than located on a single evolved Node Base station (eNB). This transmission technique forms the primary focus of the C-RAN architecture [34], which is likely to be the main enabler of CoMP operation in the near future. Network MIMO is fundamentally different to our work in that cells jointly transmit to simultaneously served UEs and thus require close coordination, as opposed to our proposed solution, where each UE only associates with a single eNB at a time and which applies reassignments to improve MU-MIMO operation.

Finally there is the concept of Dynamic Cell(/Point) Selection (DCS), in which a cell-edge UE is dynamically reassigned between serving eNBs, selecting whichever provides the highest instantaneous SINR [47]. Our work can be seen as a considerable step beyond this in which MU-MIMO pairing also forms a major deciding factor in reassignments, meaning that MU-MIMO gains are effectively taken advantage of.

To the best of our knowledge, our work in [104] was the first that investigated actively reassigning UEs from one cell to a neighbouring cell in order to increase capacity through enabling use of MU-MIMO.

In this chapter, we derive upper and lower bounds on the spectral efficiency gains experienced by the reassigned UE as a result of these reassignments taking into account the enabling of MU-MIMO transmissions and SINR losses. In the simplest case, through MU-MIMO operation in the target cell, spectral efficiency gains of up to 35% can be achieved by the reassigned UE, relative to SU-MIMO operation in the original cell.

Through extensive simulation we investigate our proposed reassignment mechanism in detailed indoor and outdoor scenarios, or more specifically, an apartment block scenario and a commercial city centre square scenario, both of which contain multiple small cell base stations and experience interference from nearby macrocell base stations. We perform proportional fair MU-MIMO user scheduling both before and after the reassignment to effectively assess what gains in spectral efficiency are obtainable through use of such a reassignment mechanism.

From detailed system level simulations we find that from the reassignment of a single UE the set of UEs for which MU-MIMO is enabled achieves spectral efficiency average gains of 0.6 (from 2.16 to 2.78) and 0.7 (from 1.43 to 2.20) bps/Hz for indoor and outdoor deployments, compared to our baseline scenario in which UEs are statically allocated to the eNB of highest Reference Signal Received Power (RSRP), which, for our investigated network configuration, is the SINR-maximising eNB allocation method. This corresponds to the enabling of MU-MIMO for the reassigned UE on 96% and 94% of scheduled subbands, respectively.

4.2 System Model

The system modelling used in this chapter makes use of many components already described in this thesis. We consider a downlink network scenario containing N_{UE} UEs, each with N_r receive antennas, and N_{eNB} base stations, each with N_t transmit antennas. The base stations constitute a heterogeneous mix of tri-sector macrocell eNBs and open subscriber group small cell eNBs all of which operate within the same

frequency band. Our UEs of interest are those served by the small cell eNBs and as such the macrocell eNBs are simply modelled as sources of inter-cell interference.

We assume the ability of having coordinated control among sets of small cells. As such, we assume the existence of a centralised coordinator which is responsible for initiating the reassignments. We don't, however, assume tight synchronisation, such as is necessary for CoMP or C-RAN.

As was discussed in Section 2.2.2, for the case of operator-deployed small cells there are a range of emerging architectures which, through centralisation of the baseband processing and simplification of the small cell devices, greatly simplify network reconfiguration and updates, and also enable differing levels of coordination between centrally-processed small cells. Current deployments offer only simple coordination while future deployments will offer more advanced capabilities. Examples include the baseband pool [35] and C-RAN [34] network architectures. As, in this work, transmissions to each UE are performed by a single small cell individually only simple coordination is required. This ability to work without tight synchronisation or continuous near-ideal information exchange between cooperating nodes presents the main advantage of this mechanism over CoMP Joint Transmission (JT).

We apply the MU-MIMO signal model which we earlier presented in Section 2.1.2. This was expressed in Eqn. (2.4) as

$$\mathbf{y}_k = \underbrace{\mathbf{H}_{k,O}\mathbf{w}_k x_k}_{\text{Desired Signal}} + \underbrace{\sum_{j=1}^{n_{MU}-1} \mathbf{H}_{k,O}\mathbf{w}_j x_j}_{\text{Multi-User Interference}} + \underbrace{\sum_{l=1}^{N_{eNB}-1} a_l \mathbf{H}_{k,l} \mathbf{W}_l \mathbf{x}_l}_{\text{Inter-Cell interference}} + \mathbf{n}_{W,k}. \quad (4.1)$$

In this chapter, we model interfering small cell eNBs as always active, meaning that $a_l = 1 \forall l$.

At the receiver end, a Minimum Mean Square Error (MMSE)-Interference Rejection Combining (IRC) filter \mathbf{g}_k is applied, as described in Section 3.2.7. As stated in Eqn. (2.6), the resulting SINR can be expressed as

$$\gamma_k = \frac{|\mathbf{g}_k \mathbf{H}_{k,O} \mathbf{w}_k|^2}{|\mathbf{g}_k \sum_{j=1}^{n_{MU}-1} \mathbf{H}_{k,O} \mathbf{w}_j|^2 + |\mathbf{g}_k \sum_{l=1}^{N_{eNB}-1} a_l \mathbf{H}_{k,l} \mathbf{W}_l|^2 + \sigma^2 \mathbf{I} \|\mathbf{g}_k\|^2}. \quad (4.2)$$

MU-MIMO Channel Quality Indicator (CQI) is modelled as in Section 3.2.4, in which UEs feed back SU-MIMO CQI which is then converted to MU-MIMO CQI at the eNB taking into account the expected level of residual unsuppressed Multi-User Interference (MUI) (Δ_{MUI}) for each number of co-scheduled UEs n_{MU} .

As specified in section 3.2.6, spectral efficiencies of MU-MIMO operation are obtained as the MU-MIMO expected throughputs per unit bandwidth scaled by the number of co-scheduled UEs, n_{MU} , thus reflecting the sharing of that resource through MU-MIMO.

As part of the reassignment mechanism investigated in this chapter UEs must estimate, both whether they experience a sufficiently high SINR in order to justify the usage of MU-MIMO in a neighbouring cell, and the magnitude of the gain in spectral efficiency they can expect from doing this. These are assessed using neighbouring cell Channel State Information (CSI). This is modelled, as discussed in Section 3.2.5, as wideband Implicit CSI under the hypotheses of MU-MIMO usage and the original cell remaining active after the reassignment occurs. Alternatively, implicit statistical information feedback (time averaged) could be used, although this is not investigated in this work as the use of wideband CQI and Precoding Matrix Indicator (PMI) feedback already provides averaging in the frequency domain.

Compared to similar multi-cell coordinated operations, the level of additional required feedback to be computed and shared is low. In coordinated scheduling, the required feedback would often consist of PMI information and expected CQI improvements (similar to our expected UE spectral efficiency improvements), for all cells within the CoMP Set of each cell-edge UE [40], where the CoMP Set is the set of cells within a cooperation area and will change depending on which cell clustering method is used. This is similar to the required feedback for our mechanism, but with the exceptions that CSI is only required for the considered UEs (cell-edge UEs unable to perform MU-MIMO in their original cell, rather than all cell-edge UEs), and that instead of requiring this for all cells within the CoMP Set we require it only for the target eNBs. Additionally, as after a reassignment transmissions can continue unaltered for a reasonable period of time, the frequency with which this feedback is required is low.

Further, if CoMP Joint Transmission were in use, subband CQI, PMI and wideband Rank Indicator (RI) would be required as often as every Transmission Time Interval (TTI), which is considerably more feedback than is used in our

proposed method.

While not further explored in this work, we expect that our proposed concept could easily be extended to operate along with coordinated scheduling. For example, if a considered UE gets reassigned to a neighbouring cell for MU-MIMO, then we expect that coordinated scheduling could be used to reduce the interference from the original cell, using the knowledge of what the precoder of the reassigned UE was before reassignment. If both methods were used together, as the CSI feedback required for *MU-MIMO across Small Cells* is a subset of that required for coordinated scheduling, no additional feedback would be needed.

Needless to say, however, the channel feedback required for our proposed MU-MIMO based reassignments does exceed that of standard DCS, which does not require neighbouring cell PMIs.

4.3 Reassignment Mechanism

As summarised in Figure 4.1 the *MU-MIMO across Small Cells* mechanism consists of four main steps:

- A. Selection of Considered UEs
- B. Checking for Target UEs
- C. Selection of UE to reassign from the set of Reassignable UEs
- D. Reassignment

The following subsections explain the steps of the mechanism in more detail, while subsection 4.3.D also discusses bounds on the spectral efficiency gains of reassigned UEs.

4.3.A Selection of Considered UEs

UEs each decide if they should be considered for MU-MIMO-based cell reassignment. UEs which are unable to benefit from MU-MIMO usage in their original cell but

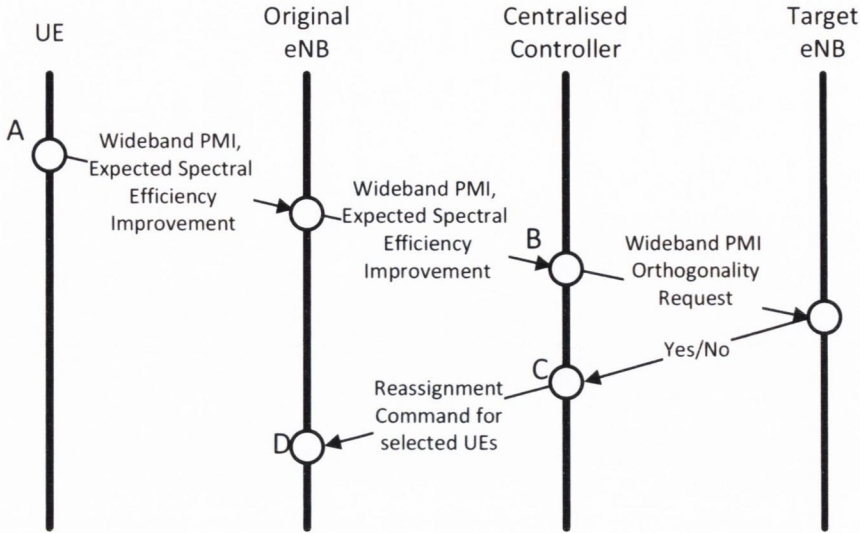


Figure 4.1: Process of MU-MIMO Across Small Cells. A. UE checks if it should be considered for reassignment. Considered UEs: Cannot be paired for MU-MIMO in their original cell, and must have a MU-MIMO spectral efficiency in the target cell higher than their SU-MIMO spectral efficiency in the original cell. B. Centralised Controller uses the wideband PMI to check if there is a UE in the target cell to pair the considered UE with. C. If there are multiple UEs which could benefit from reassignment within a given reassignment window, the centralised controller selects the reassignable UE with the largest predicted increase in spectral efficiency. D. Selected UE is reassigned.

have a sufficiently high SINR to benefit from MU-MIMO in a neighbouring cell, are considered for MU-MIMO-based cell reassignment. In other words, the considered UEs are those that meet the constraint

$$2\overline{r_{2,k,T}} > \overline{r_{1,k,O}} > 2\overline{r_{2,k,O}}I_{k,O} \quad (4.3)$$

where $\overline{r_{n_{MU},k,e}}$ denotes the expected rate averaged over all Resource Blocks (RBs). $n_{MU} = 1$ and $n_{MU} = 2$ denote SU- and two-layer MU- MIMO, respectively, and the averaged rates are scaled by the number of co-scheduled UEs (layers) to correspond to the average spectral efficiencies in bits per RB, as discussed in Section 3.2.6. $I_{k,e} \in \{0, 1\}$ indicates the presence of other UEs attached to eNB e with quantised channels semi-orthogonal to that of UE k . As before, $e = O$ regards the currently attached(/original) eNB, while $e = T$ regards a neighbouring eNB targeted for reassignment.

As previously stated, a UE cannot know prior to the scheduling process whether it will be paired for MU-MIMO in the following transmission or not. Similarly, it cannot

know the channel vectors (or PMIs) of other nearby UEs. It will, however, be able to know (from the Demodulation Reference Signal (DM-RS) in LTE Rel.10 [14] or MU-MIMO downlink power offset indicator in LTE Rel.8 [15]) if recent transmissions to it had been performed using MU-MIMO. As this work considers a slowly changing channel, the UE can expect that, if MU-MIMO was used in the previous transmission, there is another UE in the cell with a quantised channel vector semi-orthogonal to its¹. This means that $I_{k,O}$ can be specified as

$$I_{k,O} = \begin{cases} 1 & \text{if MU-MIMO was used in the} \\ & \text{previous received transmission} \\ 0 & \text{otherwise.} \end{cases} \quad (4.4)$$

Each UE considered for cell reassignment (considered UE) feeds back a single wideband PMI for each target neighbouring cell, as well as its predicted gain in spectral efficiency ($2\overline{r_{2,k,T}} - \overline{r_{1,k,O}}$), to the central coordinator. This expected spectral efficiency improvement assumes that MU-MIMO will be used in all (scheduled) RBs after reassignment, which we will see later is very close to what actually occurs.

4.3.B Check for Target UEs

In the work of this thesis we term *Target UEs* as UEs attached to the target neighbouring eNB with which the Considered UE could potentially be co-scheduled for MU-MIMO. The central controller checks for Target UEs by checking if the quantised channel vector of any UE in the neighbouring cell is semi-orthogonal to that of the considered UE (corresponding to the fed back wideband PMI), and that their SINR is not too low for MU-MIMO operation to be beneficial. Any Considered UEs for which a Target UE exists are termed *Reassignable UEs*.

¹To reiterate, a UE's quantised channel vector is obtained from its PMI as the pseudo-inverse of the PMI's corresponding codeword. These quantised channel vectors indicate the directionality of a UE's channel. In MU-MIMO, UEs with close to orthogonal (semi-orthogonal) quantised channel vectors are co-scheduled as this corresponds to low losses in SINR.

4.3.C Selection of UE to reassign from the set of Reassignable UEs

If, in any given reassignment window, multiple reassignable UEs exist, the UE with the highest expected spectral efficiency improvement is reassigned to its corresponding target eNB.

This step of selecting a single UE to reassign prevents cases where a UE gets reassigned into a target cell only to find that the UE it was supposed to pair with has also been reassigned into another cell at the same time, which could result in negative reassignment gain. This also allows us to isolate the gains of a single reassignment in our analysis. As an alternative to reassigning the UE with the highest expected spectral efficiency gain, random selection of the UE to reassign would remove the need to communicate the expected spectral efficiency improvement, although at the cost of a less favourable reassignment gain, on average.

4.3.D Reassignment

From this reassignment, gains in spectral efficiency can be achieved by both the UE that gets reassigned to the neighbouring cell (reassigned UE) and the UE that is targeted for pairing for MU-MIMO transmission (target UE).

To provide some bounds on the spectral efficiency gains of the Reassigned UE (RUE) we consider best and worst case reassignment scenarios.

As a best case we consider the case where the SINR of the RUE in the target cell is equal to its SINR in the original cell. In this case the RUE spectral efficiency increases as it would if MU-MIMO were used in the original cell. In other words, the best case fractional increase in spectral efficiency is given by

$$\text{Gain}_{\text{best case}} = \frac{2\overline{r_{2,k,T}} - \overline{r_{1,k,O}}}{\overline{r_{1,k,O}}} \quad (4.5)$$

where $\overline{r_{n_{MU},k,T}} = \overline{r_{n_{MU},k,O}}$.

As a worst case we consider the case where $2\overline{r_{2,k,T}}$ is *just* greater than $\overline{r_{1,k,O}}$ and for whatever reason (perhaps a change in the target cell PMI post reassignment)

MU-MIMO cannot be used in the target cell after reassignment. In other words the expected MU-MIMO spectral efficiency for the target cell is just greater than the original cell SU-MIMO spectral efficiency and the left constraint in (4.3) is just met, although after reassignment only SU-MIMO operation is feasible. The (negative) gain resulting from such a reassignment can then be expressed as

$$\text{Gain}_{\text{worst case}} = \frac{\overline{r_{1,k,T}} - \overline{r_{1,k,O}}}{\overline{r_{1,k,O}}} \quad (4.6)$$

where $\overline{r_{1,k,O}} = 2\overline{r_{2,k,T}}$.

These bounds are plotted in Figure 4.2 taking $\overline{r_{n_{MU},k,e}} = \log_2(1 + \text{SINR}_{n_{MU},k,e})$ where $\text{SINR}_{n_{MU},k,e}$ is the SINR of UE k , when co-scheduled as one of n_{MU} UEs, attached to eNB e . $\text{SINR}_{n_{MU},k,e}$ for different values of n_{MU} are obtained using (3.6) for the different values of unsuppressed multi-user interference at the receiver, Δ_{MUI} . The bounds assume that $\text{SINR}_{n_{MU},k,O} \geq \text{SINR}_{n_{MU},k,T} \forall n_{MU},k$ and that Single-User MIMO beamforming on a single spatial stream (SU-MIMO) is performed before reassignment. They also do not take into account SINR (/CQI) fluctuations, scheduling, Modulation and Coding Scheme (MCS) quantisation, or reference signalling overheads, which we will instead collectively take into account in Section 4.4 through extensive simulation.

We see that, for the receiver we consider in this work for which $\Delta_{MUI} = 0.05$, a maximum increase in spectral efficiency of 34.6% can be achieved when the original cell (and also target cell) SINR is roughly 11 dB, while the loss from an ineffective reassignment can be as much as 25.7% at an original cell SINR of 15.5 dB. This said, as we will see later, the likelihood of a reassignment which does not result in the enabling of MU-MIMO is very low, resulting in a positive gain in spectral efficiency for the vast majority of reassignments.

Further, if we consider the case where the SINR in the target cell may be greater than in the original cell, performing a reassignment which does not necessarily enable MU-MIMO operation can be seen as the Dynamic Cell(/Point) Selection (DCS) case, while performing a reassignment specifically targeting the enabling of MU-MIMO operation can be seen as our solution. $\text{Gain}_{\text{best case}}$ in Figure 4.2 essentially shows the gain of MU-MIMO operation over SU-MIMO if both cases operate under the same SINR conditions. Therefore, if we consider two similar reassignments, one DCS and one MU-MIMO-based, both resulting in the same post-reassignment SINR, we

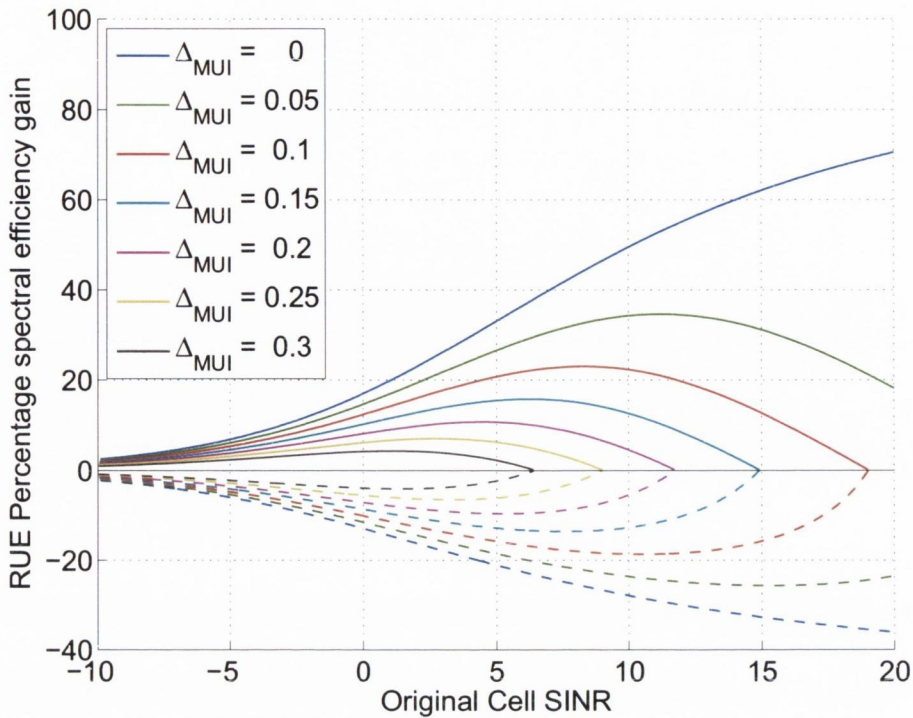


Figure 4.2: Best (solid line) and worst (dashed line) case spectral efficiency reassignment gains to the reassigned UE for different levels of unsuppressed multi-user interference at the receiver, Δ_{MUI} .

can compare these two cases using $\text{Gain}_{\text{best case}}$, where $\text{Gain}_{\text{best case}}$ demonstrates the potential gain of our MU-MIMO-based reassignments over DCS.

4.4 Simulation Results

The main focus of this section is on the assessment of the potential gains of the proposed *MU-MIMO across Small Cells* mechanism. We present these gains as the increases in spectral efficiency and MU-MIMO usage resulting from these MU-MIMO-based cell reassignments compared to a baseline in which MU-MIMO-based reassignments do not occur and in which UEs are simply assigned to the eNB of highest RSRP. In [7] it is shown that in cases where a UE can associate to eNBs from a number of different network tiers (e.g. femto, pico, micro, macro), each with different multi-antenna capabilities (numbers of antennas and set of multi-antenna transmission mode that can be used), in order to maximise all UEs' average SINRs, a tier-dependent offset to the RSRP should be applied. In our

case, as UEs can only associate to a single tier (femto/pico, depending on the scenario) in which the multi-antenna capabilities of all eNBs are the same, RSRP-based assignment *does* provide the average SINR-maximising solution, and hence was chosen as our baseline for comparison.

The following results were obtained using the Matlab-based Vienna LTE System-Level simulator [8] making use of a number of the extensions to the simulator described in Section 3.2 including MU-MIMO scheduling, MMSE-IRC reception and small cell reassignments.

Monte Carlo simulations were performed as a series of short snapshots which were averaged over a large number of iterations to generate the presented results. Once per snapshot the *MU-MIMO across Small Cells* reassignment mechanism (provided in Section 4.3) is carried out. Provided that there was an expected gain in the spectral efficiency of at least one candidate UE from reassignment, within each snapshot one UE is reassigned from its original cell to a target neighbouring cell in which it will perform MU-MIMO. Thus, we can consider the case before the reassignment as our baseline case, in which UEs are assigned to eNBs based on their RSRP, and the case after reassignment as the case which we to demonstrate, in which a MU-MIMO-based reassignment has been performed.

We investigate the performance of our mechanism in two of our implemented small cell scenarios which were introduced at the end of Section 3.3. These are the Dual Stripe indoor femtocell scenario, which is based on 3rd Generation Partnership Program (3GPP) guidelines, and the 3-picocell outdoor Munich scenario which is based on detailed pathloss computations for a real deployment location.

Each eNB has four closely spaced transmit antennas, while the UEs have two receive antennas which use MMSE-IRC. All UEs feed back a single wideband PMI for across all subbands, as described in Section 2.1.1. The PMI is selected from a codebook of possible precoders as the one which maximises the SU-MIMO mutual information summed over all subcarriers in the channel bandwidth. We consider LTE-A Rel.10 MU-MIMO operation which supports the subband use of MU-MIMO on some subcarriers and SU-MIMO on others, as well as supporting the use of Zero-forcing to orthogonalise non-orthogonal precoders in Semi-orthogonal User Selection (SUS). Despite making use of Rel.10 MU-MIMO, for 2 or 4 Transmit (Tx) antenna configurations the Rel.10 precoding codebook is the same as for Rel.8 MU-MIMO, and so the Rel.8 codebook is essentially used. As already stated, only single-layer

transmissions per UE are performed.

As the number of UEs per cell is low, we consider a maximum number of UEs that can be co-scheduled for MU-MIMO, N_{MU} , of two. This allows us to limit DM-RSs overhead, which doubles if $N_{MU} > 2$, and also provides sufficient gains to demonstrate the benefits of our proposed mechanism. Further, as our investigations in this chapter look into the gains that are achieved from the reassignment of a single UE, the most that n_{MU} can increase by as a result of any single reassignment is 1. That said, were N_{MU} increased, firstly, the reassignment algorithm designed for the base case of $N_{MU} = 2$ can be readily applied without major modification, and secondly, given that spectral efficiency increases are observed in the case where N_{MU} is limited to 2, we can expect that similar, but increased, gains would be experienced for $N_{MU} > 2$.

Both before and after the reassignments, UEs are scheduled in each TTI using the proportional fair SUS algorithm discussed in Section 3.2.2. This algorithm first selects the UE with the highest proportional fair metric to be scheduled, where the proportional fair metric is the ratio of the instantaneous achievable rate to the long-term average throughput of a UE. Next a set is computed of the UEs with wideband quantised channel vectors semi-orthogonal to the first selected UE, based on the predefined bound ϵ . From this semi-orthogonal set the UE with the highest proportional fair metric is selected to be paired with the already scheduled UE. Finally, a check is performed to ensure that the proportional fair metric of the two selected UEs both performing MU-MIMO exceeds that of the first UE alone performing SU-MIMO. Whichever of the two possibilities has the better performance is then used.

The scheduling algorithm is performed independently by each eNB and is repeated for each RB, beginning with the RB containing the UE of highest proportional fair metric and then iterating over all others. The average throughput is recomputed at each iteration. This ensures fast convergence to a proportional fair equilibrium [90].

At the start of simulation all UEs are assigned to the eNB with the highest RSRP. In the Munich scenario all 3 small cells initially contain the same number of UEs, while in the Dual Stripe scenario UEs are uniformly distributed over the entire region of interest. We assume that all UEs are slow moving and that reassignments of UEs between cells can be performed quickly. UEs can only be reassigned between small cells and cannot be reassigned to a macrocell.

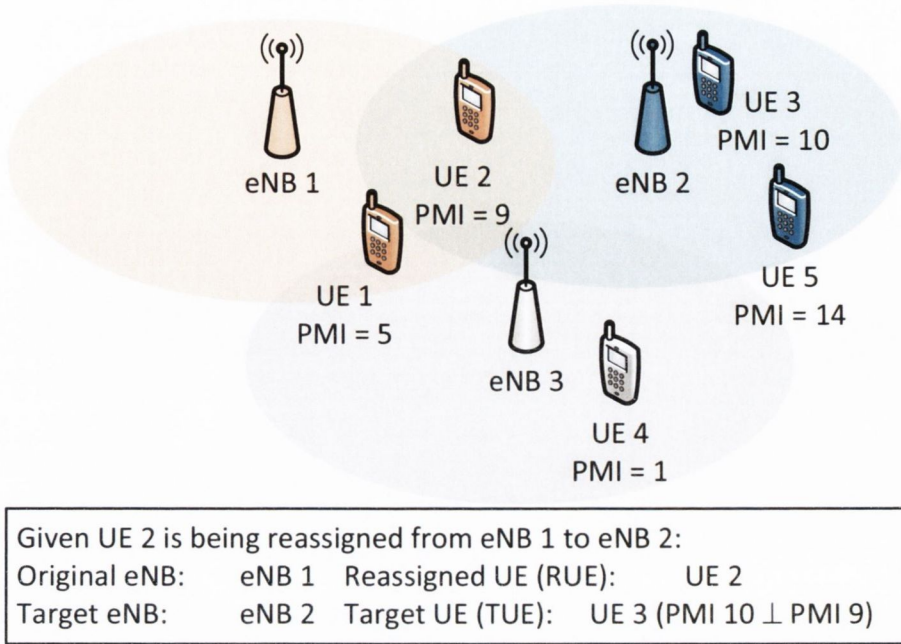


Figure 4.3: Reassignment example. UE colours indicate the Small Cell eNB to which each UE is initially attached. Considered UEs are those that cannot be paired for MU-MIMO in their original cell, although are within range of a neighbouring cell (UEs 1 & 2). Reassignable UEs are candidate UEs which have a suitable MU-MIMO pair in their target cell (UE 2 can be paired with UE 3 as their precoders are orthogonal and both have sufficiently high SINR). Of the reassignable UEs the one with the highest expected gain is chosen to be reassigned. In this case there is only one reassignable UE (UE 2).

A comprehensive list of simulation parameters used is provided in Table 4.1.

In order to review the terms used, Figure 4.3 provides an example illustration of a network in which a reassignment occurs, identifying each set of UEs discussed in the presented results. As before, to simplify explanation, this figure assumes that the PMI of each UE stays the same regardless of which eNB it is served by; in reality, the channels from different eNBs are likely to instead correspond to different PMIs.

4.4.1 Reassignability of UEs

Figure 4.4 shows the percentages of the UEs, firstly, considered for reassignment, and secondly, specified by the mechanism as reassignable, for the indoor Dual Stripe scenario with a Deployment Ratio (DR) of 0.4. The error bars on the graphs in this section represent the 95% confidence interval.

We notice from Figure 4.4 that the number of reassignable UEs does not change

Table 4.1: Simulation parameters

Scenario-specific Parameters	3GPP Dual Stripe (Indoor)	Munich 3-pico cell case (Outdoor)
Pathloss Model	3GPP Dual Stripe [95]	WiSE [102]
Fast Fading Model	Winner II [108]	Winner II [108]
Deployment Densities	Deployment Ratio = {0.2, 0.4, 0.6, 0.8}	Average inter-site distance = 57m
Number of Small Cells	{8,16,24,32} on average	3
Number of Interfering Macrocell Sectors	3	12
Initial UE distribution	Uniform across ROI	Uniform per cell
General Parameters		
Bandwidth	10 MHz	
Snapshot Length	10 TTIs (5 TTIs before and 5 TTIs after reassignment)	
Channel Feedback Delay	1 ms	
UE Speed	3 Km/h (0.83 m/s)	
eNB Antenna Configuration	4 Tx antennas Cross-polarised 0.5 λ spacing, -45 $^\circ$ /,45 $^\circ$ slants	
UE Antenna Configuration	2 Rx antennas Cross-polarised 0.5 λ spacing, 0 $^\circ$ /,90 $^\circ$ slants	
MIMO Transmission Scheme	SU-MIMO: single layer MU-MIMO: max. 2 UEs, 1 layer per UE	
Precoding Codebook	Rel.8 4 Tx codebook	
Average No. UEs per Cell	{1, ..., 8}	
Initial Cell Selection	Maximum RSRP	
Feedback (To assigned cell)	Subband CQI, wideband PMI for all UEs	
Feedback (To centralised controller for the Target neighbouring cells)	Wideband PMIs and wideband expected spectral efficiency improvement for UEs considered for reassignment.	
MU-MIMO Δ_{MUI}	0.05	
UE Scheduling	Proportional Fair SUS	
SUS const ϵ	0.1	
Traffic Model	Full Buffer	
Inter-cell Interference Model	4 Tx SU-MIMO with random PMI	
Feedback Overhead	31.15%	

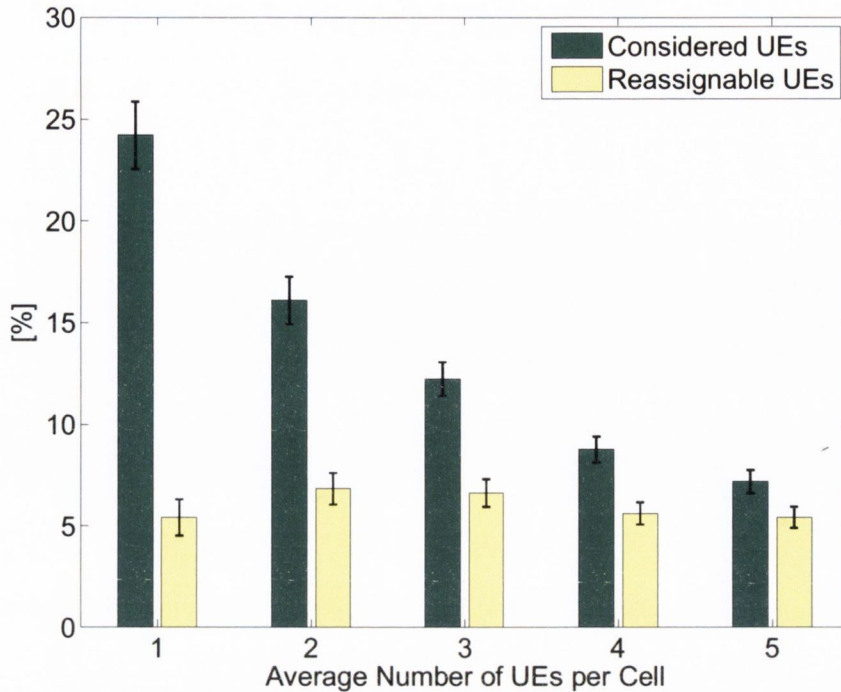


Figure 4.4: Percentage of UEs that are considered for reassignment and are decided to be reassignable in the Dual Stripe scenario with DR of 0.4.

much with the number of UEs per cell. This is because when the number UEs per cell is low, there are fewer UEs to pair with in the original cell (so the proportion considered for reassignment will be higher), although there are also fewer UEs to pair with in neighbouring cells and so the percentage for which there is benefit from reassignment remains almost the same.

Table 4.2 compares the percentages of UEs suitable for reassignment in the indoor and outdoor scenarios. The indoor scenarios are denoted by their deployment ratio $\{0.2, \dots, 0.8\}$ and the outdoor scenario is labelled Munich. We found that, in the indoor scenario, up to a certain point, increases in cell density result in more nearby cells available to be reassigned into and hence more reassignable UEs, before interference from those cells becomes a limiting factor and the percentage of reassignable UEs decreases again. Further, the percentage of reassignable UEs in the outdoor scenario is significantly higher than for the indoor scenarios (almost double). This is because in the outdoor scenario there are no walls (penetration losses) between neighbouring cells, and, even though the distance between two adjacent sites may be higher, the difference in SINR between them is often lower, leading to a higher proportion of UEs being in range of multiple cells at

Table 4.2: Cell reassignment observations

Scenario	Indoor				Outdoor
	0.2	0.4	0.6	0.8	Munich
Deployment Ratio					
Reassignable UEs [%]	5.8	6.3	6.2	5.8	11.1
Probability of MU-MIMO usage after reassignment [% of RBs]	95.6	95.9	96.6	95.1	93.5

once.

4.4.2 Increases in MU-MIMO usage due to UE reassignments

As a metric to assess the levels of MU-MIMO usage which are enabled by MU-MIMO-based reassignments we use the difference between the percentage of scheduled RBs(/subbands) in which MU-MIMO was used before reassignment, and after reassignment. Equivalently this can be termed the increase in the average number of spatial layers utilised in scheduled RBs expressed as a percentage. This is averaged over a small number of TTIs both before and after the reassignment in each snapshot. We call this the increase in MU-MIMO usage, and we assess this from the points of view of the UE that gets reassigned (RUE), the set of UEs for which the reassignment enables MU-MIMO usage increases (RUE and Target UE (TUE)), and finally for the average UE in either the cell from which the RUE originates or the target cell into which the RUE is reassigned.

Figure 4.5 illustrates these increases for a deployment ratio of 0.4. These are also separated into the before and after cases in Figure 4.6. We observe that both the RUEs and TUEs go from almost never using MU-MIMO to almost always when the number of UEs per cell is low, and, while the increase in MU-MIMO usage lessens as the UEs density grows, it still remains high.

From the simulations we found that, after reassignment, almost all transmissions to either the RUE or the TUE are performed using MU-MIMO. The percentages of RBs in which MU-MIMO is used after reassignment for each deployment ratio are shown in Table 4.2. Inevitably, a small number of RBs will remain in which MU-MIMO cannot be used after reassignment due to frequency selective fading. This means that the main reason why the increase in MU-MIMO usage is less when there are more UEs per cell is because more UEs were performing MU-MIMO transmissions

before being reassigned.

The main focus of the mechanism presented in Section 4.3 is on achieving spectral efficiency gains through enabling increased MU-MIMO usage by the UE that gets reassigned. The mechanism, however, does not take into account whether the TUE is performing MU-MIMO prior to the reassignment. This check was omitted because, even when the TUE is performing MU-MIMO before the reassignment occurs, the RUE can still obtain MU-MIMO gains by pairing with the TUE on some RBs, while the TUE continues to perform MU-MIMO with its previous pair on others. Additionally, in these cases, the RUE will often not only be capable of pairing with the TUE but also the previous pair of the TUE, further increasing RUE MU-MIMO usage probability. This is due to the structure of the precoding codebook. Hence, the choice of TUE does not depend on whether it is using MU-MIMO beforehand, so it makes sense that the probability of TUE MU-MIMO usage before the reassignment would increase with the number of UEs per cell.

In the case of the RUE it is less clear why, despite being considered for reassignment, it would have a non-zero probability of MU-MIMO usage before reassignment. While changes in the channel are slow, there is still some low probability of a UE switching from being able to perform MU-MIMO to unable in a small number of TTIs (either a fade in the wideband CQI or else the PMI changing to result in loss of precoder orthogonality between two previously paired UEs). However, this will not change with UE density.

Revisiting the reassignment mechanism, the decision of UE suitability for original cell MU-MIMO is based on, firstly, whether MU-MIMO was performed in its previous transmission, and secondly, its spectral efficiency *averaged over all RBs*. This essentially provides a wideband estimate of MU-MIMO capability which is necessary when a UE can only be active in one cell or another. This proves appropriate when the number of UEs per cell is low (which is our scenario of focus) although as the number of UEs per cell increases the number of RBs allocated to the cell-edge RUEs will decrease significantly, decreasing the effectiveness of this estimate.

At the start of the simulation UEs are assigned to eNBs based on their Reference Signal Received Power (RSRP). Unfortunately RSRP does not take into account multi-antenna gains. This means that, in cases where the beamforming gains in

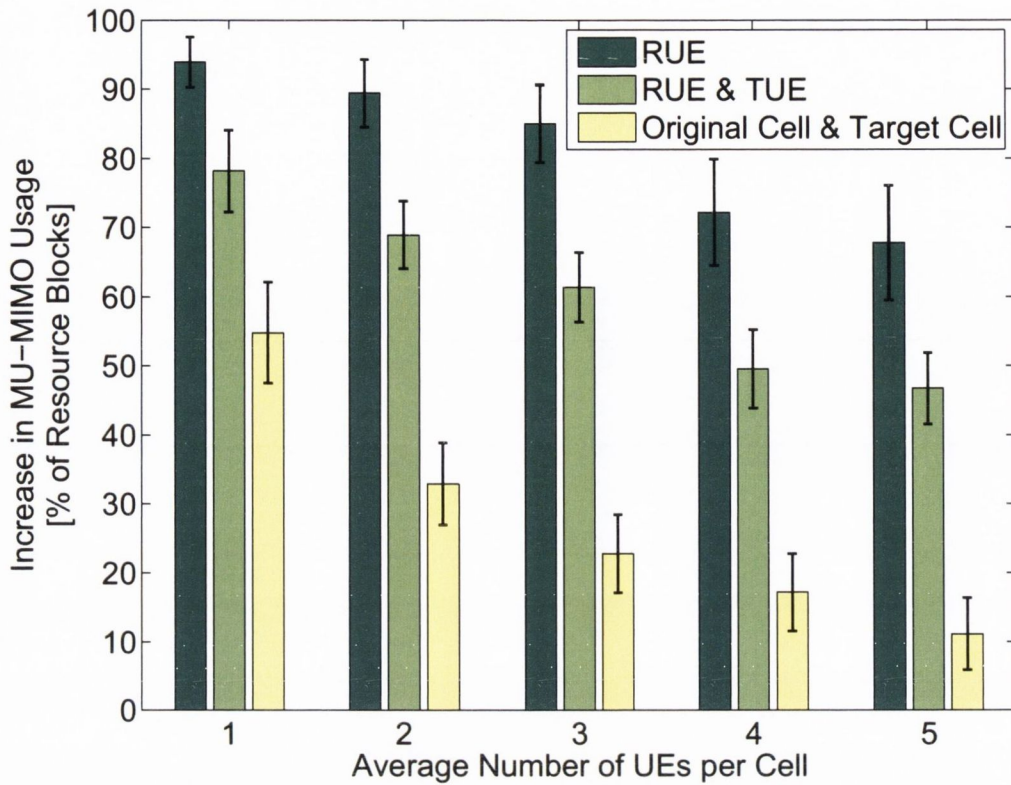


Figure 4.5: Increases in Multi-User MIMO usage resulting from reassignment for $DR = 0.4$. These increases are presented as: the increase in MU-MIMO usage by the reassigned UE alone (RUE); the average increase across the reassigned and target UEs (RUE & TUE); and the increase averaged over all UEs in both the original and target cells (Original Cell & Target Cell).

the original cell are not as good as in the target neighbouring cell, the SINR in the neighbouring cell can actually be higher. This can occur in LTE systems due to the limited codebook size and, while the probability of it is low, the likelihood of an occurrence increases with the number of UEs in the network. As this can result in a spectral efficiency gain from the reassignment that is higher than normal, it is likely that reassignable UEs for which this occurs will be selected for reassignment. This can result in cases where, before reassignment, a UE could have an SINR, averaged across the bandwidth of their original cell, which would be considered too low for MU-MIMO usage and yet still have a small number of RBs with sub-band SINRs sufficiently high, on which MU-MIMO can be used. Further, as the RUE is on the cell-edge, when there are many UEs in the cell creating competition over resources, this may make up a significant portion of their allocated RBs, which can be perceived as significant pre-reassignment MU-MIMO usage or a low post-reassignment MU-MIMO




Figure 4.6: Percentage of RBs in which MU-MIMO is performed Before and After Reassignment, as well as the difference between the two (Reassignment Gain), for DR = 0.4. These are presented for the cases of: the reassigned UE alone (RUE); the mean across the reassigned and target UEs (RUE & TUE); and the mean across all UEs in both the original and target cells (Original Cell & Target Cell).

usage increase.

Nonetheless, these occurrences are uncommon, meaning that the increase in MU-MIMO usage by the RUE is always quite high. Further, as the probability of MU-MIMO usage for both the RUE and TUE is close to 100% after reassignment, the mechanism is shown to perform the desired task well, particularly in our targeted scenario of low number of UEs per cell.

Looking at the MU-MIMO usage increases averaged over all UEs in the original and target cells we see an exponential decrease with the number of UEs per cell. Within these cells only the RUEs and TUE are significantly affected by the reassignment. When there are fewer UEs these make up a larger proportion of the total number of UEs across the two cells. Hence, the increase in average UE

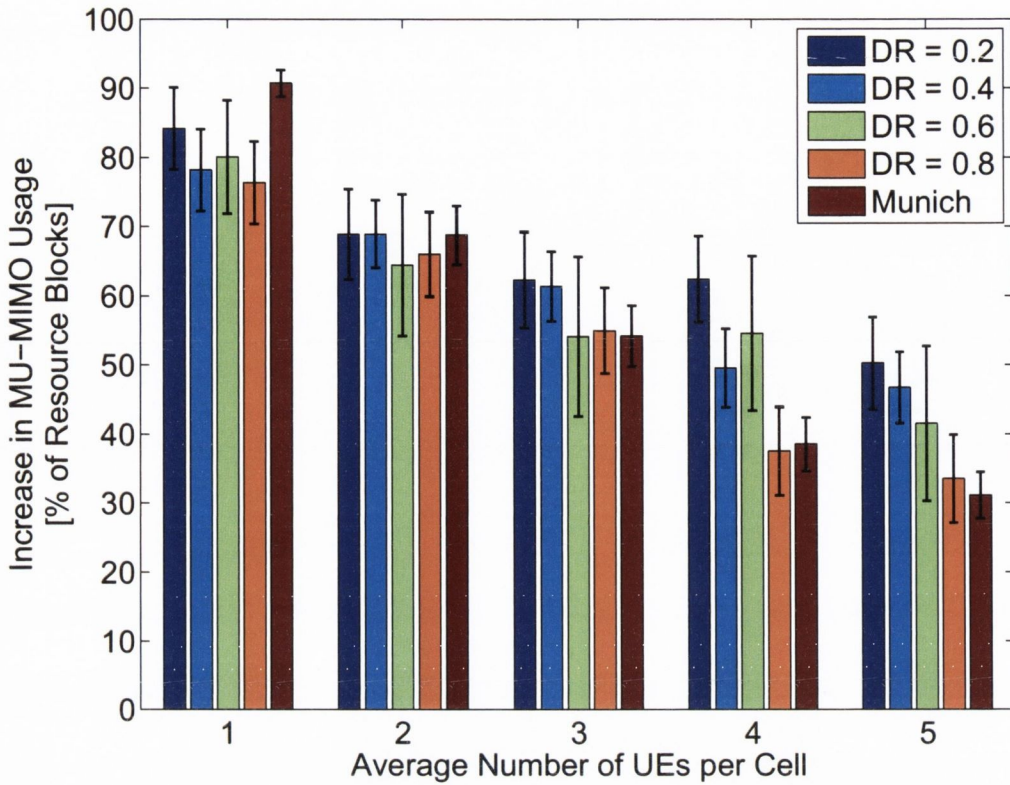


Figure 4.7: Increases in Multi-User MIMO usage resulting from reassignment for different deployment cases averaged over the RUE & TUE.

MU-MIMO usage is at its highest when the number of UEs per cell is low and gets diluted as the number of unaffected UEs grows.

Figure 4.7 demonstrates how the increases in MU-MIMO usage resulting from reassignment change with cell density. We see that, for the indoor scenarios, MU-MIMO usage increases are not affected as much by the DR as by the number of UEs per cell. Nonetheless, at low DRs the gains are slightly higher. This is because at high DRs there are more interference sources, causing more dynamic interference.

For low numbers of UEs per cell, the outdoor Munich scenario achieves a larger increase in MU-MIMO usage than the indoor scenarios. This is because the SINR difference between neighbouring cells is potentially quite small, enabling the UEs to take full advantage of any quantised channel vector orthogonalities with UEs in the neighbouring cells.

To round up, as stated earlier, it was found in [62] that, when the number of UEs per small cell is low, MU-MIMO struggles to perform well due to the lack of

appropriate UE pairings. The increases in MU-MIMO usage shown in Figure 4.7 demonstrate the ability of our proposed MU-MIMO-based reassignment mechanism to overcome the limitation experienced in [62], particularly in cases where there are initially few UEs per cell.

4.4.3 Increases in Spectral Efficiency due to UE Reassignments

Figure 4.8 presents the increases in spectral efficiency that are achieved through MU-MIMO-based reassignment for a DR of 0.4. Absolute values of the spectral efficiencies for the cases of before reassignment, after reassignment and reassignment gain are shown in Figure 4.9, where the reassignment gain is the same as the increases in spectral efficiency presented in Figure 4.8.

We start by focusing on the increases in RUE spectral efficiency, which grow with the number of UEs before levelling out. To explain this trend, we discuss three main contributing factors: the user scheduling, the total number of UEs within the cooperating region, and the initial eNB assignments.

Proportional fair scheduling allocates resources in such a way that it favours UEs with good channels but still maintains some level of fairness.

Considering the situation before reassignment, as the RUEs tend to be on the cell edge they are not particularly favoured by proportional fair scheduling. Increasing the number of UEs in the cell creates more competition over resources, meaning that before reassignment the RUE spectral efficiencies will be lower when there are more UEs per cell. On the other hand, for the TUEs, as they are, on average, closer to the cell centre, proportional fairness tends to increase their spectral efficiencies with the number of UEs per cell.

After reassignment both the RUE and TUE operate in MU-MIMO mode with high probability. This means that the RUE will be co-scheduled in the RBs of the TUE and vice versa. Both UEs will experience a gain in spectral efficiency as a result of using MU-MIMO and the RUE also receives an additional gain in spectral efficiency as a result of the TUE essentially competing for resources on its behalf.

The combination of the before and after reassignment situations regarding scheduling results in increases in spectral efficiency with the number of UEs per cell before saturation, as observed.




Figure 4.8: Increases in Spectral Efficiency resulting from reassignment for $DR = 0.4$. These increases are presented as: the increase in spectral efficiency to the reassigned UE alone (RUE); the average increase across the reassigned and target UEs (RUE & TUE); and the increase averaged over all UEs in both the original and target cells (Original Cell & Target Cell)

The second cause for the trend in the increases in RUE spectral efficiency is the *total number of UEs* served by the cooperating eNBs. When the total number of UEs increases there are more reassignable UEs to select from. This choice provides an increased reassignment gain when the network size is larger.

Thirdly, as discussed in the previous subsection, it can occur that the *initial RSRP-based eNB assignment* results in some UEs achieving a better multi-antenna gain in the target cell than the original cell. In these cases, if an MU-MIMO-based reassignment occurs, the gains in spectral efficiency from reassignment can be greater than previously anticipated. The likelihood of this also increases with the network size.

Steering focus to the averaged increase in spectral efficiency across both the RUE & TUE in Figure 4.8, in general, the gain to the RUE exceeds the gain to the TUE.

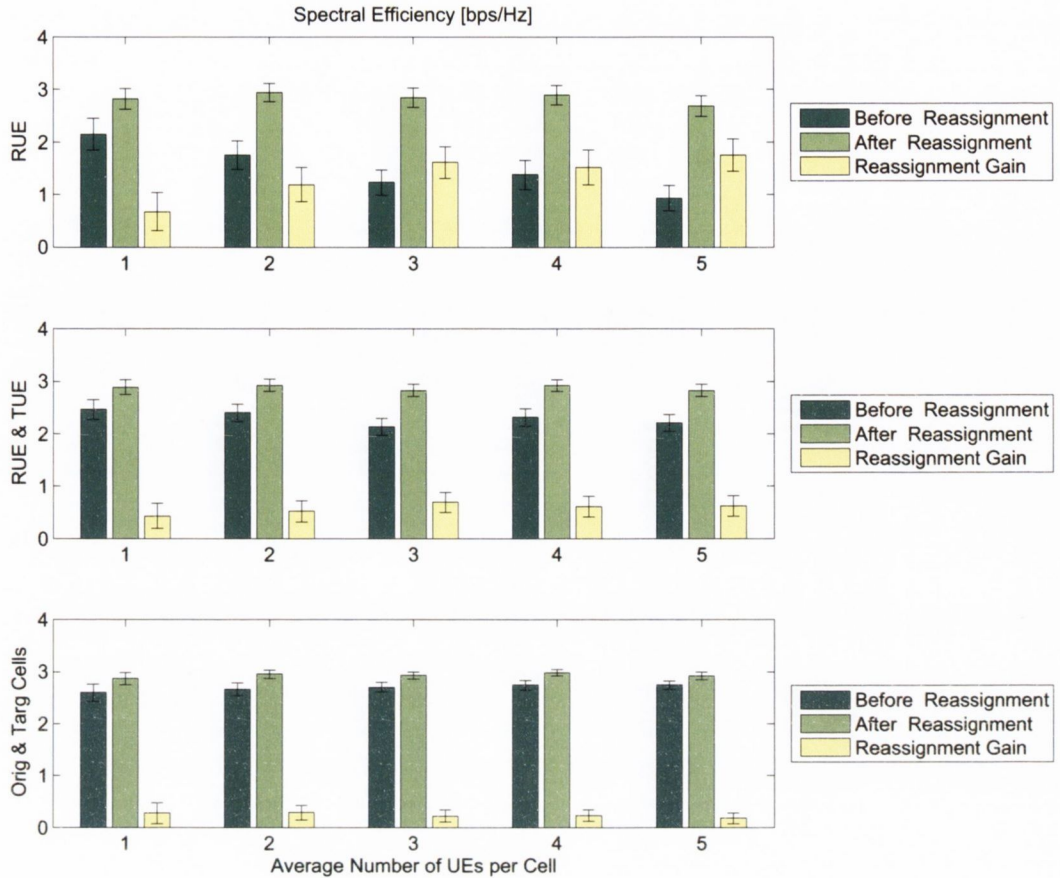


Figure 4.9: Spectral Efficiencies Before and After Reassignment, as well as the difference between the two (Reassignment Gain), for DR = 0.4. These are presented for the cases of: the spectral efficiency of the reassigned UE alone (RUE); the mean spectral efficiency across reassigned and target UEs (RUE & TUE); and mean spectral efficiency across all UEs in both the original and target cells (Original Cell & Target Cell)

Again as discussed in the previous subsection, this is because the mechanism focuses primarily on reassignment gains of the RUE. In this, whether the TUE is performing MU-MIMO prior to the reassignment is not considered and, as a result, TUEs don't achieve the same increases in MU-MIMO usability as RUEs, which results in them not achieving as high spectral efficiency gains either. Nonetheless, when a TUE does transition from fully unpaired beforehand to fully paired after reassignment (e.g. when there are few UEs per cell), its gain in spectral efficiency is often higher than that of the RUE. The reason for this is that UEs within range of only one cell observe less interference from neighbouring cells than those in range of multiple (generally RUEs), and as a result they operate at a higher SINR for which the gain from using MU-MIMO is often greater.




Figure 4.10: Increases in Spectral Efficiency resulting from reassignment, averaged over the RUE & TUE, for different deployment cases.

While the RUE bars in Figure 4.8 (which form our main target in this work) can be seen as similar to cell-edge spectral efficiency gains, we must also ensure that the average UE isn't detrimentally affected by the reassignment operation. For this, we also look at the average increase in spectral efficiency across all UEs in either the original cell or target cell. Even though, as already discussed, the reassignment gains are diluted by larger numbers of UEs per cell, the average UE benefits from the reassignment in all simulated cases.

Figure 4.10 shows the increases in spectral efficiency averaged over the RUE and TUE for the considered indoor and outdoor scenarios. The corresponding absolute values of the (RUE and TUE) spectral efficiencies for the cases of before reassignment, after reassignment and reassignment gain are shown in Figure 4.11, where the reassignment gain is the same as the increases in spectral efficiency presented in Figure 4.10. For the same reasons as just discussed increasing DR results in improved gains but faster saturation. Further, the outdoor scenario,

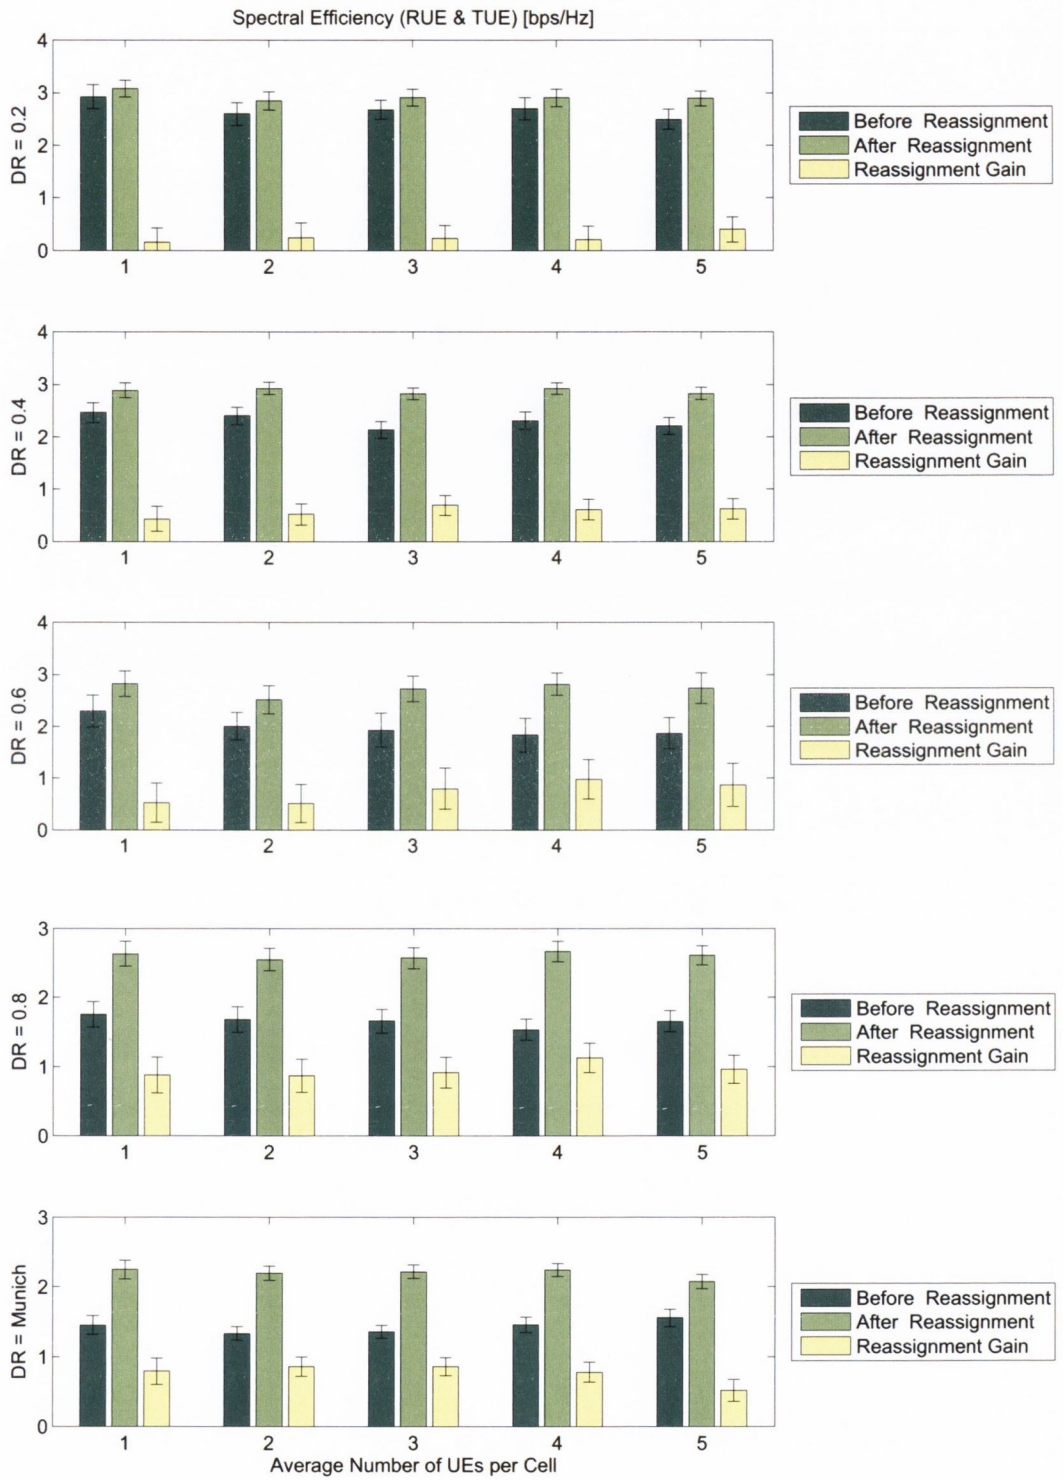


Figure 4.11: Spectral Efficiencies, averaged over the RUE & TUE, Before and After Reassignment for different deployment cases.

despite having significantly fewer cooperating small cells, achieves comparable reassignment gains to the indoor scenario with a DR of 0.8 at low numbers of UEs per cell. This implies that the *MU-MIMO Across Small Cells* mechanism is particularly suited to outdoor scenarios in which there are no penetration losses from walls between cooperating small cells.

4.4.4 Time Scale Considerations

The mean length of time for which the channel of a UE (and hence its PMI) can be assumed to be stationary is given by its coherence time, T_c , and is inversely proportional to the velocity of the UE [18]:

$$T_c = \frac{c_0}{2f_c v} \quad (4.7)$$

where c_0 is the speed of light ($3 \times 10^8 m/s$), f_c is the carrier frequency ($2.1 \times 10^9 Hz$) and v is the velocity of the UE in $[m/s]$.

The total time taken to perform a standard handover from one cell to another is specified by 3GPP as $D_{handover}$ and is computed as the sum of the Radio Resource Control (RRC) procedure delay and the interrupt time $T_{interrupt}$ [109, Section 5.1]. More specifically, this is the time between when the downlink handover command is received and when the UE recommences uplink transmissions. The RRC procedure delay is the time taken for the RRC to be re-established and is given in [110, Section 11.2] as not to exceed $15ms$, while $T_{interrupt}$ is given in [109, Section 5.1] as not exceeding $50ms$ given that the target cell is known (as is the case in our scenario). Therefore the total handover time $D_{handover}$ is $\leq 65ms$.

This means that in order for the channel coherence time not to exceed the handover time the speed must be kept below $1.1m/s$ ($4km/h$). While this is above the speeds considered in small cell scenarios, which are designed to satisfy low mobility UEs and as such consider UE velocities ranging from 0 to $3km/h$.

Next we consider, for our case, how regularly reassignments are likely to be required. This again is determined by the channel coherence time although as the coherence time represents the mean time for the channel to change and does not give much guide to the variance, it is hard to say how performance will be degraded

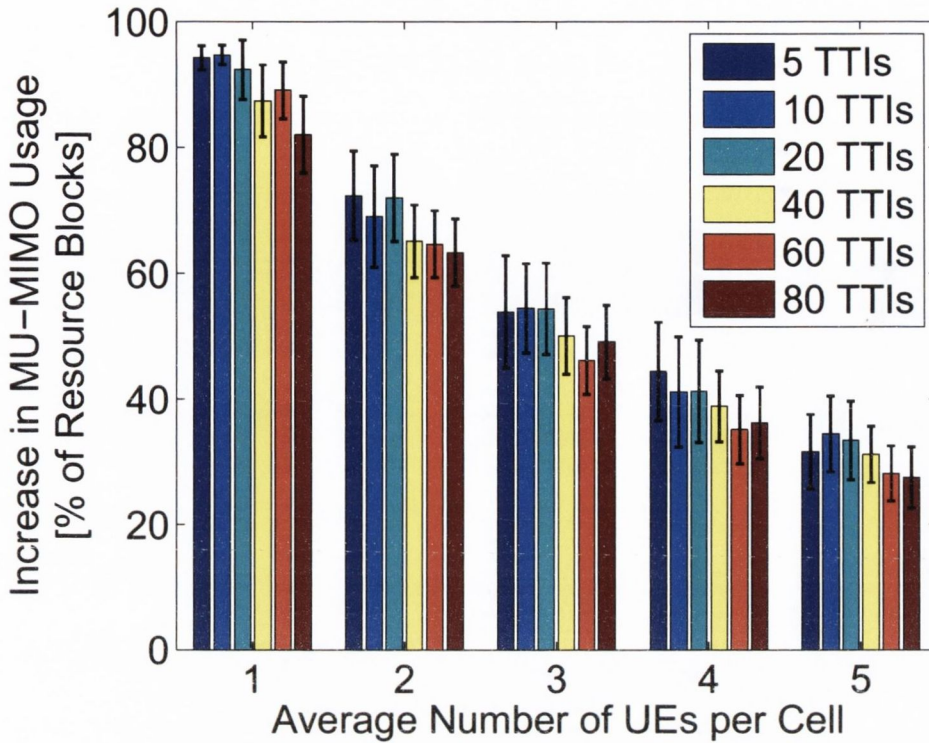


Figure 4.12: Multi-User MIMO usage resulting from reassignment for different periods of time before and after reassignment. Munich outdoor deployment scenario, RUE & TUE case.

as the time between reassignments is increased. We investigate this through simulation by varying the snapshot window size over which the increase in MU-MIMO usage is averaged. This is shown in Figure 4.12 where, as before, reassignments of UEs in the simulations are fast (by which, specifically, we mean that the reassignment is performed within 1 TTI). The figure shows that increasing the time between reassignments from 5 TTIs to 80 TTIs results in roughly 13% less of an increase in averaged MU-MIMO usage levels.

Within LTE, coordinated schemes are generally grouped into dynamic ($\geq 1ms$), semi-static ($\geq 1s$) and static (days), where all CoMP schemes are performed dynamically. As the timescales we consider are between tens and hundreds of milli-seconds, dynamic coordination for our reassignments is required, although some flexibility can be allowed for.

4.5 Summary

In this chapter, we have proposed and investigated a new method of increasing MU-MIMO gains in small cell networks, which we call *MU-MIMO across Small Cells*. This method involves the coordinated reassignment of UEs which were unable to be co-scheduled for MU-MIMO in their original cell, to neighbouring cells in which they can perform MU-MIMO.

Two small cell scenarios were investigated: an indoor apartment block scenario, and an outdoor scenario corresponding to a commercial square in the city centre of Munich, Germany. Both were assessed from the perspectives of spectral efficiency and MU-MIMO usability. It was found that in the outdoor scenario, more than 11% of the UEs served by the three small cells in the scenario met the reassignment criteria, which was almost double what we saw for the indoor scenario. We found that all the investigated cases achieved positive increases in both spectral efficiency and MU-MIMO usage across the affected cells as a result of the reassignment. We also found that after the reassignment, MU-MIMO is used in upwards of 95% of all RBs scheduled to either the reassigned or target UE.

In the following chapter this work will be extended to also take energy efficiency considerations into account. MU-MIMO-based reassignments of UE will be performed with the added consideration that if a cell becomes empty it can be switched to a sleep state in order to conserve energy. As a result we modify the reassignment mechanism toward an objective of maximising the number of small cell eNBs which are switched to sleep modes. This is subject to a minimum expected RUE spectral efficiency gain for any performed MU-MIMO-based reassignments.

5 MU-MIMO Across Small Cells: Energy and Spectral Efficiency Increases

This chapter is largely based on the results and contributions of our paper entitled "*Energy and Spectral Efficiency gains from Multi-User MIMO-based Small Cell Reassignments*" [111] which was submitted to *IEEE Global Communications Conference (GlobeCom) 2015*. However additional insights are also included.

5.1 Introduction

As we showed in the previous chapter, by selectively reassigning User Equipments (UEs) between neighbouring small cells in such a way as to enable Multi-User MIMO (MU-MIMO) usage where it was previously unavailable, we can obtain significant increases in UE spectral efficiency. The work of this chapter combines this concept with the use of small cell sleep modes to obtain simultaneous increases in both spectral *and* energy efficiency. To enhance this operation the reassignment mechanism of the previous chapter was appropriately refined.

To recap briefly; traditionally, use of MU-MIMO has been largely limited to macrocell scenarios. The reason for this is that, in order to keep the interference between two simultaneously served UEs (Multi-User Interference (MUI)) to a minimum, the co-scheduled UEs must have sufficiently uncorrelated (semi-orthogonal) channels. In small cell scenarios, where the number of UEs per

cell is small, it is not always possible to find suitable UE sets. However, in small cell scenarios a UE can often be in range of multiple small cells. By reassigning UEs between neighbouring evolved Node Base stations (eNBs) in such a way as to increase the number of suitable UE pairings, additional MU-MIMO spatial multiplexing gains can be enabled [104]. This can be seen as similar to enlarging the search space for suitable UE MU-MIMO pairings.

Given this scenario, once a UE is reassigned to a target neighbouring cell, if the host cell is no longer serving any active UEs, it can enter an idle state where radio transmissions are temporarily suspended, achieving high energy savings. This additionally results in a reduction in interference and pilot pollution problems, which are often observed in dense deployments. Moreover, considering that electricity costs account for 20-30 percent of network operational expenses [112], there is no lack of motivation for operators to become more energy efficient. Within this context, this chapter introduces a novel scheme where UEs are selected for reassignment to neighbouring eNBs in order to save energy by reducing the number of active small cell eNBs serving UEs, while simultaneously achieving spectral efficiency gains by enabling use of additional spatial layers through MU-MIMO.

The use of small cell sleep states has been proposed as a solution to limit energy consumption increases as small cell numbers grow large. This is of particular importance considering the rate at which small cell deployments are being rolled out, with the total worldwide expected to reach 100 million by 2020 [113].

As discussed in [112], sleep mode techniques often fall into one of three categories, differentiated by the method used for their reactivation. These are: Small Cell Controlled, in which an RF sniffer is utilised to identify potential UEs to serve [114]; Core Network Controlled, in which, as in our work, small cells are put into sleep states and re-awoken by a centralised core network element; and finally UE controlled, in which UEs emit periodic wake-up signals for surrounding small cells.

Traditionally these works have taken advantage of low traffic conditions and the energy saving possibilities they present, although more recently, given the density of small cell deployments and the associated cell redundancies/over provisioning, the work has evolved to also include user association considerations [65, 115–117]. From this perspective, one work related to ours is [115], in which centralised decisions are made to switch a portion of the active small cell eNBs to sleep modes, from

which the resulting gains in energy efficiency and Signal to Interference and Noise Ratio (SINR) (as a result of reductions in inter-cell interference) are studied through system level simulation. Our work further improves the UE performance by selectively deactivating small cell eNBs in such a way that achieves both MU-MIMO spatial multiplexing and interference reduction spectral efficiency gains.

[65] looks at uplink Cloud Radio Access Networks (C-RANs) with centralised joint decoding of the received signals of multiple multi-antenna small cell base stations (this is the uplink equivalent of Coordinated Multi-Point (CoMP) Joint Transmission (JT)). The small cells compress their received signals so that they can be transported over a finite capacity backhaul to the central processor. The authors formulate the compression into a rate maximising optimisation problem, into which they also include a sparsity-inducing term (a base station activity cost) creating a joint optimisation problem in which the number of active small cells is also minimised. Different from [65] we do not consider joint encoding or decoding amongst cooperating base stations. Instead MU-MIMO transmissions are performed separately by each small cell base station, which considerably reduces the computation and coordination overheads.

As in the previous chapter's reassignment process, a set of UEs which can benefit from MU-MIMO-based reassignments is constructed. If all reassignable UEs are reassigned at once a game of cat and mouse can ensue where no suitable new pairs are created due to different halves of the proposed UE pair being reassigned in opposite directions. For this reason, in this work we select the set of reassignable UEs which maximises the number of small cells which are emptied of attached UEs as a result of the reassignment. These small cell eNBs with no more UEs to serve can then be switched to a sleep/idle state and, in doing so, their energy consumption can be dramatically decreased. An example of this is illustrated in Figure 5.1. The fact that small cell scenarios consider low mobility UEs means that the frequency with which reassignments will need to be made is not excessive.

As a result, this solution can achieve simultaneous gains in spectral efficiency and energy efficiency, compared to current operation in which UE to small cell assignments are based on Reference Signal Received Power (RSRP) [6]. We perform system level simulations to quantify the achievable gains in two realistic scenarios: one indoor and one outdoor, where the indoor scenario is the same as in the previous chapter while the outdoor scenario is an enlarged version (21 picocells) of the that used earlier. In

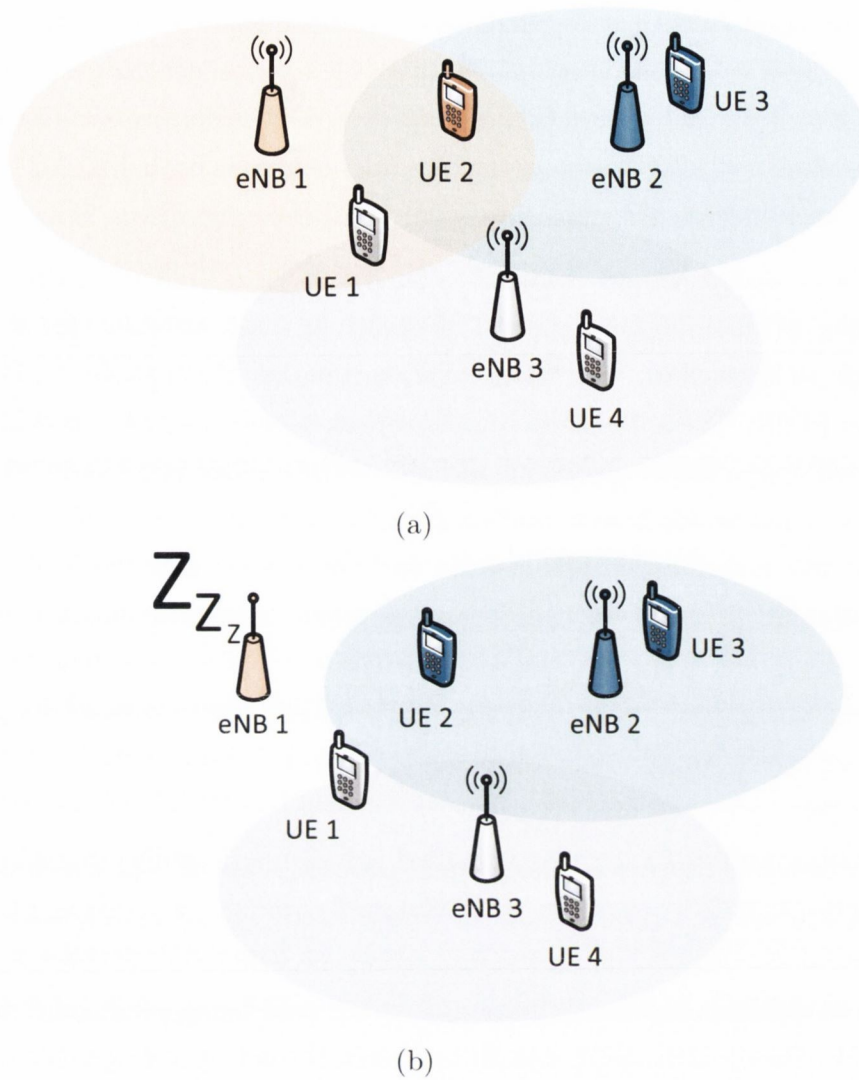


Figure 5.1: UE colours indicate the Small Cell eNBs to which each UE is attached. (a) Consider that by reassigning UE 2 to eNB 2 MU-MIMO transmissions simultaneously serving UEs 2 & 3 can be enabled (e.g. if UE 2 is within range of eNB 2 and the channels of UEs 2 & 3 are semi-orthogonal). (b) Additionally, by reassigning UE 2 to eNB 2, as there are no more UEs attached to eNB 1, eNB 1 can be switched to a sleep state to conserve energy.

In addition to being able to achieve both of these gains simultaneously, we investigate the possibility of trading-off small amounts of UE spectral efficiency for additional gains in energy efficiency.

5.2 System Model

The system model used in this chapter is similar to that used in the previous chapter; however with a few extensions.

As in the previous chapter our network contains N_{UE} UEs, each with N_r receive antennas, and N_{eNB} base stations, each with N_t transmit antennas. The base stations constitute a heterogeneous mix of tri-sector macrocell eNBs and open subscriber group small cell eNBs all of which operate within the same frequency band. Our UEs of interest are those served by the small cell eNBs and as such the macrocell eNBs are simply modelled as sources of inter-cell interference.

Again, the small cell eNBs are governed by a central coordinator which is responsible for initiating reassignments, and which, in this chapter, also specifies when small cell sleep states are to be used. We use $a_e = 1$ to indicate that eNB e is *active* and $a_e = 0$ to indicate that eNB e is in an *idle/sleep* state. As in the previous chapter this results in an MU-MIMO SINR which can be expressed as

$$\gamma_k = \frac{|\mathbf{g}_k \mathbf{H}_{k,O} \mathbf{w}_k|^2}{|\mathbf{g}_k \sum_{j=1}^{n_{MU}-1} \mathbf{H}_{k,O} \mathbf{w}_j|^2 + |\mathbf{g}_k \sum_{l=1}^{N_{eNB}-1} a_l \mathbf{H}_{k,l} \mathbf{W}_l|^2 + \sigma^2 \mathbf{I} \|\mathbf{g}_k\|^2}. \quad (5.1)$$

Channel State Information (CSI) estimation is also modelled similarly to the previous chapter; however, in this work reassignments are only performed if they enable the deactivation of the original cell; hence, we do not include interference from the original eNB in the target cell Channel Quality Indicator (CQI). This provides a more accurate estimate of the post-reassignment rate. Nonetheless, it is not possible to know, prior to the reassignment decision, which other neighbouring eNBs will also be deactivated. If additional neighbouring eNBs are deactivated the levels of interference will be reduced further, meaning that the target cell CQI provides a conservative estimate of the post-reassignment rate. Compared to the previous chapter this can be summarised in the 3GPP terminology of [37] as the use of Implicit Channel State Feedback, subject to the two hypotheses of MU-MIMO usage and original cell deactivation, as opposed to the hypotheses of MU-MIMO usage and the original cell remaining active.




Figure 5.2: MU-MIMO Across Small Cells with sleep mode usage. A. UEs assess whether they should be considered for reassignment. B. The Centralised Controller checks the target eNBs for potential Target UEs with which the Considered UEs can be co-scheduled for MU-MIMO. If such Target UEs exist, the corresponding considered UEs are termed reassignable. C. The Centralised Controller selects a subset of the Reassignable UEs to reassign. In this paper, they are selected so as to minimise the required number of active eNBs. D. The selected UEs get reassigned. E. eNBs with no remaining attached UEs switch to an idle state to conserve energy.

To model small cell energy consumption we make use of the model proposed by the FP7 Energy Aware Radio and neTwork tecHnologies (EARTH) project in [74] shown earlier in Eq. (2.7).

$$P_{in} = \begin{cases} N_t(P_0 + \Delta_P P_{out}), & ACTIVE \\ N_t P_{sleep}, & IDLE \end{cases} \quad (5.2)$$

This models the small cell base station input power in Watts, P_{in} , as a function of its state (Active/Idle) as well as the transmit power used P_{out} . From this it is possible to quantify the reductions in Carbon Dioxide (CO_2) emissions and electricity costs that result from these small cell deactivations [118]; however, these metrics are not pursued in this work.

5.3 Reassignment Mechanism

The *MU-MIMO across Small Cells* reassignment mechanism of this work follows a common structure to the previous chapter (Section 4.3), augmented by the deactivation of idle small cells, and consists of five main steps:

- A. Selection of Considered UEs
- B. Check for Target UEs
- C. Selection of UEs to reassign from the Reassignable Set
- D. Reassignment
- E. Deactivation of emptied small cells.

Additionally, as necessitated by the energy consumption minimisation objective of the reassignment, the third step, in which UEs are selected for reassignment, is much changed. These steps are summarised in Figure 5.2 and elaborated on in the following subsections.

5.3.A Selection of Considered UEs

In order to avoid unnecessary exchange of feedback information, at the start of the reassignment process, UEs assess whether they should be considered for MU-MIMO-based cell reassignment.

If a UE is capable of MU-MIMO operation in its current cell it is unlikely to improve its MU-MIMO capabilities by reassigning to a neighbouring cell and so UEs which are already performing MU-MIMO are removed from the set of considered UEs. We specify these as the UEs meeting the constraint

$$2\overline{r_{2,k,O}}I_{k,O} > \overline{r_{1,k,O}}. \quad (5.3)$$

where $\overline{r_{n,k,e}}$ denotes the expected instantaneous rate averaged over all Resource Blocks (RBs). $n = 1$ and $n = 2$ denote SU- and two-layer MU- MIMO, respectively, and the

averaged rates are scaled by the number of co-scheduled UEs (layers) to correspond to the average spectral efficiencies in bits per RB. $I_{k,e}$ indicates the presence of another UE attached to eNB e with a quantised channel vector semi-orthogonal to that of UE k .

However, as previously stated, a UE cannot know prior to the scheduling process whether it will be co-scheduled for MU-MIMO in the following transmission or not. It will, however, be able to know (from the Demodulation Reference Signal (DM-RS) in Long Term Evolution (LTE) Rel.10 [14] or MU-MIMO downlink power offset indicator in LTE Rel.8 [15]) if recent transmissions to it had been performed using MU-MIMO. As this work considers a slowly changing channel, the UE can expect that, if MU-MIMO was used in the previous transmission, there is another UE in the cell with a quantised channel vector semi-orthogonal to its, meaning that $I_{k,O}$ can be specified as

$$I_{k,O} = \begin{cases} 1 & \text{if MU-MIMO was used in the} \\ & \text{previous received transmission} \\ 0 & \text{otherwise.} \end{cases} \quad (5.4)$$

Of the UEs incapable of MU-MIMO usage in their current eNB, the ones considered for reassignment are those which can expect to benefit from use of MU-MIMO if attached to a neighbouring eNB by a certain amount:

$$2\overline{r_{2,k,T}} - \overline{r_{1,k,O}} > \tau. \quad (5.5)$$

where τ is an offset indicating how much higher the target cell expected rate must be, which can be either positive or negative, and $\overline{r_{2,k,T}}$ is obtained from neighbouring eNB reference signals as outlined in Section 3.2.5.

The parameter τ has been introduced in this chapter to provide control over the tradeoff between spectral and energy efficiency. A positive value of τ indicates a strict requirement on the minimum expected Reassigned UE (RUE) spectral efficiency gains, while a negative value of τ indicates the willingness to sacrifice some spectral efficiency for a reduction in energy consumption.

Then, these UEs which cannot perform MU-MIMO operation in their original

cell, and meet the neighbouring cell MU-MIMO rate constraint, are labelled the *Considered UEs*. These Considered UEs then share a single wideband Precoding Matrix Indicator (PMI), for each target neighbouring eNBs, with the central coordinator via their currently attached eNBs.

5.3.B Check for Target UEs

This step is unchanged from the previous chapter. The central controller checks for Target UEs by checking if the quantised channel vector of any UE in the neighbouring cell is semi-orthogonal to that of the considered UE (corresponding to the fed back wideband neighbouring cell PMI), and that their SINR is not too low for MU-MIMO operation to be beneficial. Any Considered UEs for which a Target UE exists are termed *Reassignable UEs*.

5.3.C Selection of UEs to reassign from the Reassignable Set

The choice of UEs to reassign can determine the number of eNBs which can be deactivated (put into a sleep mode) as a result of the reassignment, as well as the increases in spectral efficiency and MU-MIMO usage that can be achieved.

In this work we formulate the decision of which UEs to reassign as a set covering problem in which we select UEs to reassign so as to *maximise the number of deactivated eNBs*. In this, the set of all UEs $k \in \mathcal{U}$ must be covered (served) by the set of active base stations. However a UE k can only be served by an eNB $e \in \mathcal{E}$ if *either* e is the original eNB of k *or* k can be reassigned to e .

This set covering problem can be expressed as follows:

$$\begin{aligned} & \text{minimize} && \sum_{e \in \mathcal{E}} (a_e) \\ & \text{subject to} && \sum_{e \in f(k)} (a_e) \geq 1 \quad \forall k \in \mathcal{U} \\ & && a_e \in \{0, 1\} \quad \forall e \in \mathcal{E} \end{aligned}$$

where a_e indicates that eNB e is active and $f(k)$ indicates the set of eNBs which can serve UE k .

In this work we consider two methods of solving this problem. In order to assess the performance of the mechanism under optimal reassignment selection (maximum number of deactivated eNBs) we solve this problem using the IBM ILOG CPLEX Optimiser. Despite our problem size being relatively small and CPLEX finding solutions in a reasonable amount of time, the set covering problem is NP-hard, meaning that for real implementations a heuristic solution would be required. For this reason we also assess the performance of the reassignment mechanism when Chvatal's well-known greedy heuristic set covering solution [119] is applied. This algorithm proceeds by initially assuming all base stations to be idle and then recursively activating the base station which can serve the largest number of unserved UEs. This is repeated until all UEs are covered by at least one base station. This heuristic solution achieves an approximation ratio of $\Theta(\ln N_{UE})$ whilst operating in polynomial time.

5.3.D Reassignment

All UEs attached to any of the small cell eNBs selected for deactivation are reassigned to their respective target eNBs.

5.3.E Deactivation of emptied small cells

Once a small cell base-station has finished serving a UE or when a UE is reassigned to a neighbouring cell, the small cell base station can enter the idle mode regime to minimise energy consumption and to avoid creating unwanted interference. However, there is a need for effective mechanisms that can detect the presence of a new UE and subsequently wake up the base station. Conceptually, this can be done either in a distributed fashion and by the base station itself, or centrally and with the aid of a separate controller.

The simplest way to realise the distributed wake up mechanism is to equip the eNB with an RF sniffer which constantly monitors the detected energy in the uplink band. Once this quantity exceeds a certain threshold, the base station can interpret this as a sign indicating the presence of a UE in the vicinity of the small cell [114]. The issue with this method is that setting the detection threshold optimally is difficult

and that, contrary to the scenario in [114] where the UE communicates with an underlying macrocell eNB, in dense deployment scenarios, the UE is more likely to be served by the neighbouring small cell base station, which essentially implies that the UE transmits at lower power, making it more difficult to detect.

In the centralised approach, however, a separate network element triggers the small cell base station to wake up. This has the advantage that the base station will not be triggered unnecessarily if the UE can be served by other existing eNBs. The challenging task is to determine when a UE falls under the coverage area of a given small cell base station. This can be done by using RF-fingerprints as a metric to estimate the location of the UE or through other advanced geolocation methods.

As our network already contains a centralised controller the centralised approach presents the best fit to our scenario. Furthermore, this approach provides a higher energy efficiency gain than the distributed approach, as less infrastructure is required to detect when a reactivation is required [112].

5.4 Simulation Results

In this section we assess what energy saving and spectral efficiency gains are achievable through use of the described reassignment and small cell deactivation mechanism. We take as a baseline for comparison the current standard practice case in which UEs are initially assigned to the eNB of strongest RSRP [6, Section 5.2.3], which, as discussed at the start of Section 4.4, for cases where all eNBs operate within the same frequency band and have the same multi-antenna capabilities, represents the SINR-maximising assignment policy. This UE association method is commonly used in commercial deployments and as such we are comparing the performance of our proposed scheme to the current status quo. We investigate how these are affected by the number of UEs per cell and the reassignment parameter τ , as well as the algorithm used to select UEs to reassign from the set of Reassignable UEs.

The simulation results are based on two scenarios: one indoor and one outdoor. For the indoor scenario we consider the 3rd Generation Partnership Program (3GPP) Dual Stripe indoor femtocell scenario described in Section 3.3.1, while for the outdoor scenario we consider the 21 picocell Munich city centre scenario described in Section

3.3.2. This outdoor scenario covers a wider area than the one investigated in the previous chapter and contains many more small cell base stations (21 as opposed to 3). This was chosen as it allows for a more averaged result in term of the number of small cell base stations which are deactivated.

Within each simulation snapshot, provided that there exists at least one empty cell in which all UEs are reassignable, UEs will be reassigned from their original eNBs to target neighbouring eNBs in which they will perform MU-MIMO and, in doing so, enable the emptied eNBs to be switched to an idle state in order to conserve energy.

The parameters used in the energy efficiency model (P_{max} , P_0 , Δ_P and P_{Sleep}) were obtained from [74], although adjusted for the 4 transmit antenna case such that the maximum total radiated power is kept consistent. These values are based on 2010 hardware capabilities, which provide a conservative estimate of the energy savings that can be obtained, as small cell sleep state capabilities have developed over time. Further, as identified in [112] the energy savings provided by centralised control of sleep modes can be in the order of 92% which is considerably higher than achieved by the 2010 technology. It is assumed that eNBs can switch to a sleep state quickly.

For consistency the majority of the remaining parameters remain the same as were used in the previous chapter. A full breakdown of simulation parameters is provided in Table 5.1. As before, error bars in all figures in this section correspond to the 95% confidence interval, and unless otherwise stated, in order to assess the reassignment mechanism performance under optimal reassignment selection (in terms of number of deactivated small cells), figures are based on the CPLEX set covering solution rather than the greedy heuristic.

5.4.1 Energy Efficiency Gains

Figures 5.3a and 5.3b show the percentage of small cell eNBs which are switched to a sleep mode as a result of the reassignment mechanism for the indoor Dual Stripe and outdoor Munich scenarios, respectively.

As can be seen in both figures the percentage of small cell eNBs switched off decreases as the reassignment metric threshold τ increases. Higher τ corresponds to a stricter constraint on the foreseen RUE reassignment spectral efficiency gains while

Table 5.1: Simulation parameters

Scenario-specific Parameters	3GPP Dual Stripe (Indoor)	Munich 21 picocell (Outdoor)
Pathloss Model	3GPP Dual Stripe [95]	WiSE [102]
Fast Fading Model	Winner II [108]	Winner II [108]
Deployment Densities	DR = 0.2	Mean inter-site distance = 37m
No. Small Cells	8 on average	21
Small Cell max Tx Power	20 dBm	24 dBm
No. Macrocell Interferers	3	18
Macrocell max Tx Power	46 dBm	43 dBm
Initial UE Distribution	Uniform per cell	Uniform per cell
P_{max}	0.025 W	0.065 W
P_0	2.4 W	3.4 W
Δ_P	4.0	8.0
P_{sleep}	1.45 W	2.15 W
General Parameters		
Bandwidth	10 MHz	
Snapshot Length	10 TTIs (5 TTIs before and 5 TTIs after reassignment)	
Channel Feedback Delay	1 ms	
eNB Antenna Configuration	4 Tx antennas Cross-polarised 0.5λ spacing, $-45^\circ/45^\circ$ slants	
UE Antenna Configuration	2 Rx antennas Cross-polarised 0.5λ spacing, $0^\circ/90^\circ$ slants	
MIMO Transmission Scheme	SU-MIMO: single layer MU-MIMO: $N_{MU}=2$, 1 layer per UE	
Precoding Codebook	Rel.8 4 Tx codebook	
Average UEs per Cell	$\{1, \dots, 4\}$	
Initial Cell Selection	Maximum RSRP	
Feedback (To assigned cell)	Subband CQI, wideband PMI for all UEs	
Feedback (To centralised controller)	Wideband PMI for each Target Neighbouring eNB	
MU-MIMO Δ_{MUI}	0.05	
UE Scheduling	Proportional Fair SUS	
SUS const ϵ	0.1	
Traffic Model	Full Buffer	
Inter-cell Interference Model	4 Transmit (Tx) SU-MIMO with random PMI	
Feedback Overhead	31.15%	





Figure 5.3: Percentage of eNBs that can be switched to an IDLE state as a result of MU-MIMO based cell reassignments for different reassignment thresholds τ in (a) the indoor Dual Stripe scenario with DR of 0.2 (total 8 femto eNBs on average), and (b) the outdoor Munich scenario (total 21 pico eNBs).

increased leniency in τ results in a larger set of reassignable UEs, allowing more small cell eNBs to be switched off.

While not shown, as τ continues to decrease, a point will be reached (roughly $\tau = -6$) where UE reassignability becomes exclusively a function of UE channel orthogonality and no longer of the neighbouring cell SINRs. In this case a UE may be reassigned to any neighbouring eNB regardless of how far apart they may be, potentially resulting in dramatic decreases in spectral efficiency and inability to use MU-MIMO in the target neighbouring cell. Instead, in the figures we show a reasonable range of τ values for which this does not occur.

The more UEs there are per eNB the less likely *all* UEs in the cell will be reassignable. Further, if there are more UEs in the original cell it is more likely that suitable MU-MIMO pairs will already exist, reducing the number of considered UEs. This results in very low probabilities of a reassignment occurring for 4 UEs per cell, even for low τ .

Comparing figures 5.3a and 5.3b we see that the proportion of eNBs deactivated is generally higher in the outdoor scenario. Figure 5.4 presents the percentage savings in small cell eNB power consumption for the case of $\tau = -1.5$. We see that the power savings in the outdoor scenario are higher for low UEs per cell, while when there are more UEs per cell the power savings become comparable.

From the power consumption parameters in Table 5.1 alone, if we consider the case of a single small cell, we can compute that the percentage power saved from sleep state usage over active mode is slightly higher in the outdoor pico case (45%) than for the femtos (indoor, 42%).

We see that in the outdoor case with 1 initial UE per cell a 15% reduction in the consumed power can be achieved. This corresponds to a 45% energy saving by the 37% percent of pico eNBs which got deactivated, or an average saving of 4.5W per pico cell. In the indoor case with the same parameters on average 2.5W per femto cell can be saved. Reducing τ further can achieve higher energy savings, although at the cost of lower (or negative) spectral efficiency gains.

Up until this point energy savings have only been presented for the optimal reassignment selection solution (optimal in terms of number of deactivated small cell eNBs for the given set of Reassignable UEs). In order to assess the performance difference between the optimal solution of the set covering problem and the most

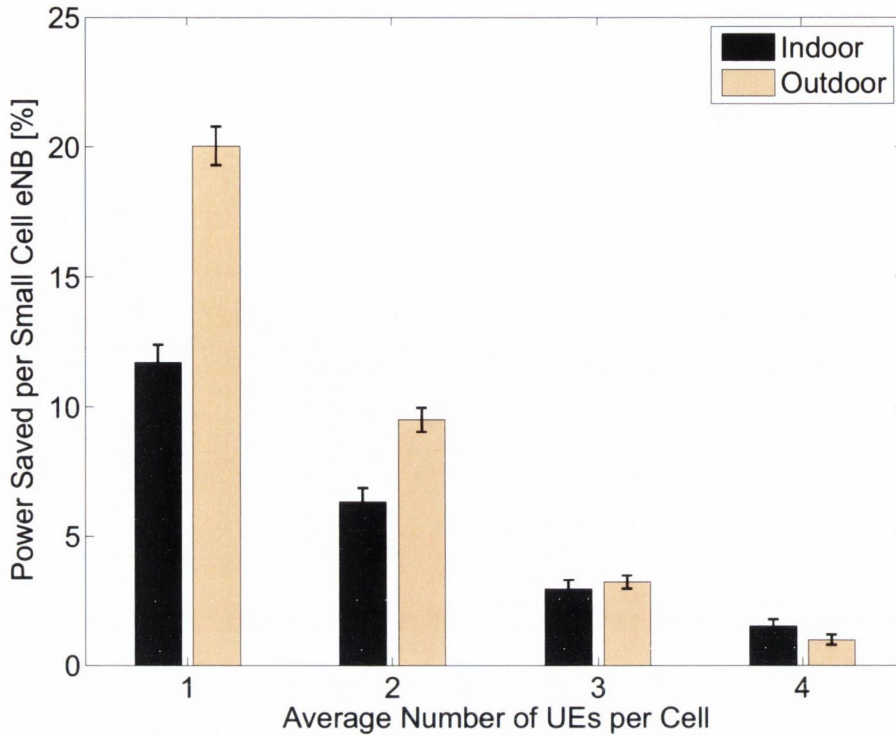


Figure 5.4: Percentage of average eNB power consumption that can be saved through switching eNBs to IDLE state as a result of MU-MIMO-based cell reassignments. $\tau = -1.5$. As can be computed through the combination of Table 5.1 and Eqn. 5.2 the power *ACTIVE* and *IDLE* state powers of an eNB in the Indoor scenario are 10 and 5.8 W, respectively, while the *ACTIVE* and *IDLE* state powers of an eNB in the Outdoor scenario are 15.68 and 8.6 W, respectively.

well-known greedy approximation [119] within our scenario Figure 5.5 compares the percentage of deactivated eNBs in the outdoor case with 1 UE initially present in each cell. As pointed out earlier the advantage to the heuristic algorithm is that it achieves polynomial time operation, while achieving the optimal solution is NP-hard and only possible for small problem sizes and loose time constraints.

The optimal solution results in, on average, 1.8 times as many deactivated small cell eNBs as the heuristic algorithm. As expected, the heuristic performs well above the approximation ratio, which provides a lower bound on performance equal to $\ln(21)+1$ for our case, where 21 is the number of UEs present in this scenario. Further, as τ increases, the performance of the greedy solution gets marginally closer to the performance of the optimal solution. This is because as τ gets larger the number of reassignable UEs decreases, hence the problem size gets smaller, corresponding to a decrease in the number of suboptimal solutions.




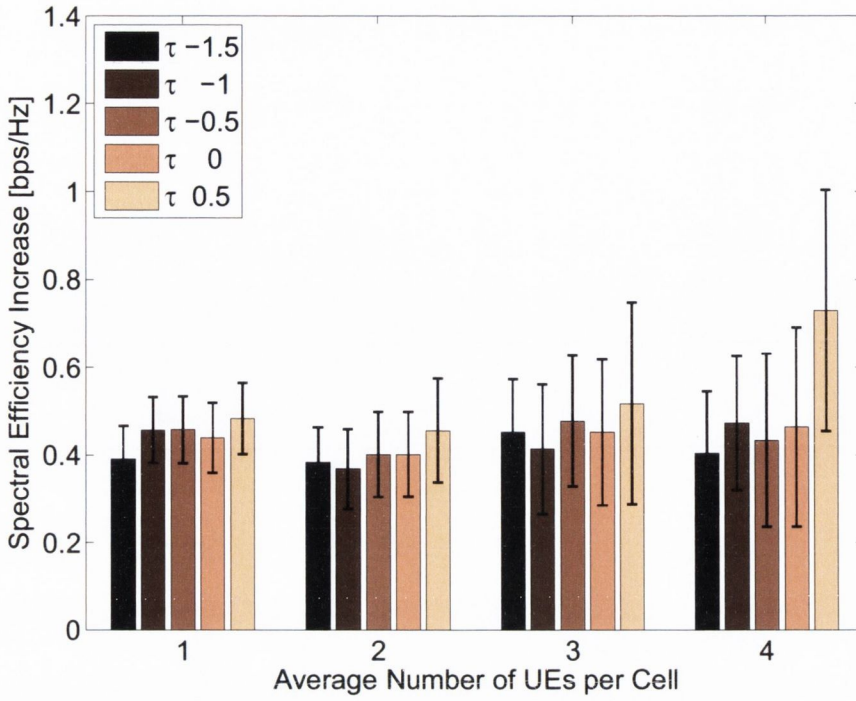
Figure 5.5: Percentage of eNBs that can be switched to IDLE as a result of MU-MIMO-based cell reassignments for both the Optimal and Greedy heuristic set covering solutions in the Munich outdoor scenario (total 21 pico eNBs) with 1 UE per eNB.

5.4.2 Increases in Spectral Efficiency


Figures 5.6a and 5.6b show the gains in spectral efficiency averaged over the UE that gets reassigned and the Target UE with which the Reassigned UE is simultaneously served for MU-MIMO in the target neighbouring cell. Equivalently, these are the gains in spectral efficiency of the enabled MU-MIMO usage over the pre-reassignment Single-User MIMO (SU-MIMO) usage, and are a function of the original and target eNB SINRs, before and after reassignment.

To provide a more complete understanding Cumulative Distribution Function (CDF) plots of the reassigned UE and target UE Spectral Efficiencies, Before and After Reassignment, are shown in Figures 5.7a and 5.7b.

As expected, a higher reassignment threshold τ generally corresponds to a higher spectral efficiency, although this is not necessarily true in all cases, and the difference is generally not much. As earlier stated, due to it not being possible to know prior to reassignment which neighbouring cells will be deactivated, τ cannot take into account

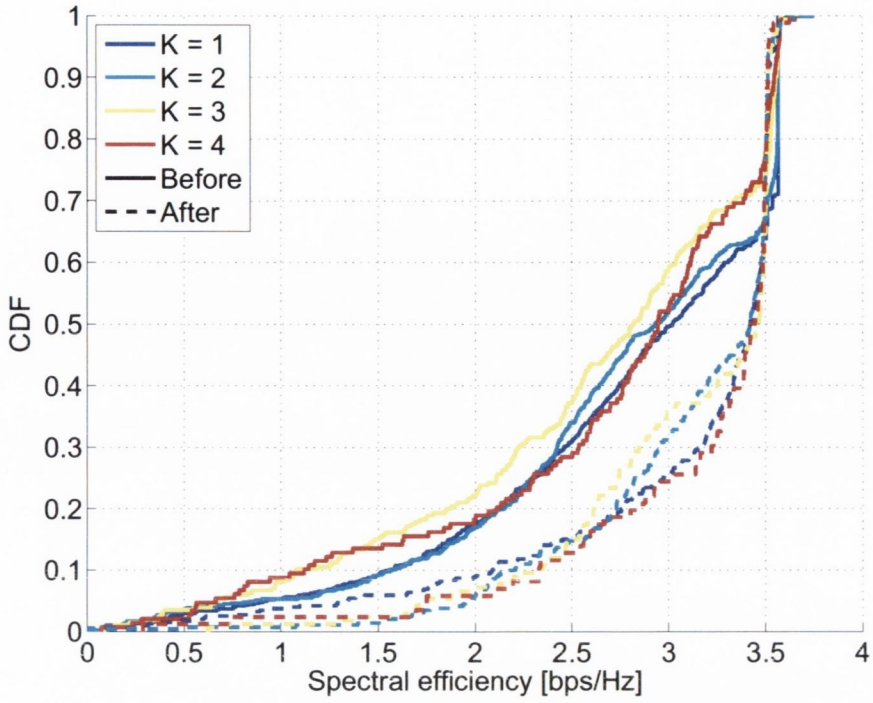


(a)




(b)

Figure 5.6: Increases in Spectral Efficiency resulting from reassignment, averaged across the reassigned and target UEs for different reassignment thresholds τ in (a) the indoor Dual Stripe scenario with DR of 0.2, and (b) the outdoor Munich scenario.



(a)



(b)

Figure 5.7: CDF of reassigned UE and target UE Spectral Efficiencies Before and After Reassignment, for reassignment thresholds $\tau=-1.5$ and number of UEs per cell K varying from 1 to 4. (a) Indoor Dual Stripe scenario with DR of 0.2, (b) outdoor Munich scenario.

gains in SINR from a neighbouring small cell eNB switching into a sleep state (thus removing the inter-cell interference it causes). Further, lower values of τ mean that more reassignments can be performed, allowing more small cells to be deactivated, and resulting in more reductions in interference. This results in the extent to which the spectral efficiency is affected by τ being less than previously anticipated. This also explains why the increases in spectral efficiency tend to exceed the values of τ .

We see that the gains in spectral efficiency in the outdoor scenario are higher than in the indoor scenario. This is due to the lack of walls between neighbouring small cell eNBs on the same street. In the indoor scenario there are walls between all eNBs, meaning that the minimum difference in SINR between two neighbouring cells is higher. This further explains why the proportion of eNBs that could be deactivated was higher in the outdoor scenario in Figure 5.3b, as this results in a larger set of reassignable UEs.

Clearly, the expected tradeoff between spectral efficiency (averaged over the reassigned and target UEs) and energy efficiency for varying τ is much less prominent than anticipated. Nonetheless, the spectral efficiency tends to increase for all values of τ , meaning that the value of τ providing the highest gain in energy efficiency can be selected without adversely affecting much the spectral efficiency of the reassigned and target UEs.

5.4.3 Discussion of Other Potential Baselines

As stated earlier, in this investigation we take as a baseline for comparison the case where all UEs are assigned to the eNB from which they receive the highest RSRP, which corresponds to the SINR-maximising allocation for our investigated scenario. That said, a range of other association policies exist within the literature, the objectives of which vary between macrocell offload, energy efficiency, bit-rate maximisation, Closed Subscriber Group (CSG) operation, and handover minimisation. In Table 5.2 we summarise these approaches and give pros and cons for each when considering our investigated scenario.

In order to bias UEs in favour of being assigned to small cells in heterogeneous networks many works consider the use of small cell range expansion offsetting to the RSRP [7, 120, 121]. The objective of this is generally to offload traffic from the

macrocell eNB and hence achieve a better distribution of load across the network; however, in [7] the authors show that, in multi-tiered networks where the different tiers (femto, pico, micro, macro) have different multi-antenna capabilities, it is also necessary to apply appropriately-chosen offsets to the RSRP-values in order to come to an average SINR-maximising UE allocation. In the scenarios we consider UEs can only be assigned to a single small cell tier in which all eNB have the same multi-antenna capabilities and operate within the same frequency band. As a result, for these association regimes, the same tier-based offset would be applied to the RSRPs of all eNBs, making it equivalent to RSRP-based assignment alone, which we apply as our comparison baseline.

From the perspective of energy efficiency, as already discussed, a number of other works consider the application of sleep states in cases of changeable [112, 114] or full-buffer [115, 117] traffic. Compared to our proposed MU-MIMO-based reassignment mechanism the full-buffer sleep mode scenarios (which are most applicable for comparison to our scenario which also considers full-buffer traffic), as they do not *require* MU-MIMO to be enabled for a cell-emptying reassignment to occur, result in a different set of UEs reassignments to our proposed solution. Hence, we expect that these approaches combined with our proposed MU-MIMO-based reassignments would achieve maximal energy efficiency without loss in spectral efficiency. Detailed investigation of this combined approach is a subject of future work.

Also from the perspective of energy efficiency another approach is proposed in [121] in which a cognitive heterogeneous network is considered in which not all small cells have a direct backhaul connection and some rely on a multi-hop relay backhaul instead. In this, by taking into account the required number of hops in the cell assignment a more energy efficient allocation can be obtained. As our scenario assumes that all small cells have a direct backhaul connection this additional consideration is not applicable to our scenario.

As bit-rate maximisation (load balancing), CSG-operation, and handover minimisation, despite being common within the small cell (and heterogeneous network) literature, do not focus centrally on energy efficiency or SINR-maximising assignment, they are presented within Table 5.2 although not discussed further.

Table 5.2: Association Approaches for Small Cells with Multiple Antennas

Association Approach	References	Pros/Cons
RSRP	[5, 7, 58, 63, 114, 120, 121]	<ul style="list-style-type: none"> + Simplicity of implementation. + Maximises the average SINR (provided that all small cells operate at the same frequency and have the same multi-antenna capabilities) [7]. - Does not explicitly take into account multi-antenna capabilities (Chapter 4). - Proposed MU-MIMO-based reassignments achieve higher spectral and energy efficiency as was shown in Figures 5.3 to 5.7.
Single-antenna Sleep Mode Techniques (Full-buffer Traffic)	[115, 117]	<ul style="list-style-type: none"> + More energy efficient than RSRP-based allocation. + Does not <i>require</i> MU-MIMO to be enabled for reassignment and, as such, results in a different set of UEs reassignments to our proposed solution. + Combined with our proposed MU-MIMO-based reassignments would achieve maximal energy efficiency without loss in spectral efficiency.
Load Balancing	[7, 120]	<ul style="list-style-type: none"> + Can achieve higher bitrates by ensuring resources are spread evenly. - Opposite to energy efficient solutions in which, to free up small cells to deactivate, the load is concentrated to as few cells as possible.
CSG	[61, 62, 115]	<ul style="list-style-type: none"> - Must conform to predefined (often suboptimal) allocation.
Power Control	[122, 123]	<ul style="list-style-type: none"> + Can provide energy efficient outcomes by consuming less transmit power and producing less interference. - Does not result in small cells making use of sleep modes, which result much more significant energy savings in small cell networks [75, 124].
Multi-Element Antenna (MEA) Antenna Selection	[58]	<ul style="list-style-type: none"> + Reduces power leakage from CSG femto cells. + Reduces number of unwanted handovers between macro and femto cell. - Target objective not energy efficiency.
JT-CoMP	[65]	<ul style="list-style-type: none"> - Need for tight synchronisation. - Multiple small cell eNBs must be active. - Mostly envisioned for Macrocell and C-RAN systems.

5.5 Summary

In this chapter we demonstrated the combined use of MU-MIMO-based UE reassignments and centralised control of small cell sleep states to achieve simultaneous increases in spectral efficiency and reductions in energy consumption.

We compared the achievable gains in two small cell scenarios, one indoor residential apartment block scenario (Dual Stripe) and one touristic outdoor city centre scenario (Stachus Square, Munich) relative to a baseline case corresponding to current common practice (RSRP-based UE-to-eNB assignments). We found that it was possible to switch in excess of 25% (indoor) and 35% (outdoor) of small cells to a sleep state whilst still achieving considerable gains in the spectral efficiency of the set of UEs for which MU-MIMO is enabled (the RUE and Target UE (TUE)). Based on the energy consumption model used [74] these correspond to power savings of 12% and 15% respectively; however, as this model is based on 2010 small cell technology, and as the energy savings through use of sleep states have increased since that time, these gains based on current or future technology will be even higher.

6 Conclusions and Future Work

6.1 Contributions and Findings

The use of multiple antennas has shown huge potential to dramatically increase the spectral efficiency of transmissions in cellular networks. In parallel, the rollout of self-configuring small cell infrastructure is leading to rapid improvements in spectrum spatial reuse. These rollouts aim to meet the high-capacity application demands expected from low mobility users in coming years. However, the use of multiple antennas to transmit over multiple spatial layers (spatial multiplexing), which is the primary multi-antenna method employed to boost capacity in macrocell environments, has not yet been properly realised in small cell environments. Spatial multiplexing gains are generally obtained in either of two ways: by transmitting over multiple spatial layers on a single link between a base station and a User Equipment (UE), or by simultaneously transmitting to multiple UEs, each on different spatial layers, in what is known as Multi-User MIMO (MU-MIMO). The use of spatial multiplexing to a single UE is often inhibited by the channel characteristics of small cell environments and by the fact that it requires multiple widely-spaced antennas at both the small cell and UE sides. At the same time, while MU-MIMO can get around these problems, its performance is also limited in small cell environments by the low numbers of UEs present in each cell.

In this thesis we have proposed and assessed two techniques for the reassignment of UEs between neighbouring small cells which overcome the MU-MIMO performance limitations experienced as a result of there being few UEs present in each cell in small cell scenarios. These techniques, which we collectively called *MU-MIMO Across Small Cells*, focused on spectral efficiency and combined energy and spectral efficiency gains,

respectively. In the course of this work a number of original contributions and findings were made. In the following we recap over these contributions and findings, before discussing possible extensions to these in future work.

In Chapter 2 we provided an overview of the background and state of the art on the topics of advanced multiple-antenna techniques and small cell deployments, specifically focused on areas of interest to this work. We summarised how Single-User MIMO (SU-MIMO), MU-MIMO and Coordinated Multi-Point (CoMP) techniques operate and the extent to which they have been included into Long Term Evolution (LTE) standardisation, as well as comparing the functionality of common CoMP techniques to the subject matter of our work. We reviewed the literature on multi-antenna usage techniques in small cell networks, in which there are strong focuses on coverage control and required deployment density considerations. We looked briefly at a number of small cell architectures focused on enabling coordination in such scenarios, and finally we discussed a number of efforts to reduce small cell energy consumption further. As well as the discussions of related work provided within Chapter 2, chapter-specific related work comparisons were provided in the introductions of each subsequent chapter.

Chapter 3 contained two primary focuses. The first was on qualitative and quantitative comparison of a representative set of potential system and network level simulators for Extremely Dense Network (EDN) investigations and the provision of guidelines for other researchers in the selection of the appropriate simulator for their desired scenario. This took into perspective key considerations for EDN simulations such as link level abstraction, modelling of interference, Radio Resource Management, different traffic types, backhaul and upper layer procedures, simulator ease-of-use, structure and modularity, supported deployment scenarios, and scalability for networks containing large numbers of nodes. Our second focus in this chapter was on the simulator extensions which we performed as part of this work to allow us investigate the proposed MU-MIMO-based reassignment mechanisms. These included system-level MU-MIMO implementation, related MU-MIMO scheduler, feedback and spectral efficiency implementations, as well as the inclusion of indoor and outdoor small cell deployment scenarios.

In Chapter 4 we proposed and assessed the performance of the first of our two investigated coordinated MU-MIMO-based reassignment mechanisms. The goal of this mechanism is to reassign UEs between adjacent cells to arrive at an assignment

in which the mutual channel orthogonality between UEs in the same cell is increased, hence enabling the use of additional spatial layers for MU-MIMO transmissions. The performance of the reassignment mechanism was assessed via Monte Carlo system level simulation in two small cell scenarios: a residential indoor apartment block scenario, and an outdoor scenario corresponding to a commercial square in the city centre of Munich, Germany. In each snapshot, the UE in the network with the highest expected spectral efficiency gain from MU-MIMO-enabling reassignment was selected to be to reassigned. It was found that in the outdoor scenario, more than 11% of the UEs served by the three small cells in the scenario met the reassignment criteria, which was almost double what we saw for the indoor scenario. We found that all the investigated cases achieved positive increases in both spectral efficiency and MU-MIMO usage across the affected cells as a result of the reassignment compared to our baseline case in which UEs are assigned to evolved Node Base stations (eNBs) based on Reference Signal Received Power (RSRP) alone (no reassignments). We also found that after the reassignment, MU-MIMO is used in upwards of 95% of all Resource Blocks (RBs) scheduled to either the reassigned or target UE.

In Chapter 5 we demonstrated the combined use of MU-MIMO-based UE reassignments and centralised control of small cell sleep states to achieve simultaneous increases in spectral efficiency of the Reassigned UE (RUE) and Target UE (TUE) and reductions in energy consumption. The reassignment mechanism presented in this chapter was posed as a set covering problem in which we want to minimise the number of small cell base stations required to cover (serve) the set of UEs in the network. In this, a small cell base station can only "cover" a UE if either it is the current serving base station of that UE (the base station providing the highest received power), or if the UE can obtain a minimum gain in spectral efficiency as a result of an MU-MIMO-based reassignment into that cell. Following this, all UEs originally assigned to an emptiable small cell were reassigned to neighbouring small cells, at which point the original cell was switched to a sleep state. This resulted in energy savings, increased MU-MIMO usage, and reductions in inter-cell interference compared to our baseline case in which, again, UEs are assigned to eNBs based on RSRP alone, representing current practice. As in the previous chapter, we investigated this in both indoor and outdoor (residential and urban) small cell deployment scenarios. We compared the use of a well-known polynomial-time heuristic solution to the set covering problem to the optimal

solution in our investigated scenario. It was found that the polynomial-time heuristic achieved 56% of the number of deactivated small cells of the NP-hard optimal solution. We found that, using optimal reassignment, it was possible to switch in excess of 25% (indoor) and 35% (outdoor) of small cells to a sleep state whilst still achieving considerable gains in the RUE and TUE spectral efficiency. These corresponded to power savings of 12% and 15%, respectively, based on the energy consumption model proposed in [74], and as the efficiency savings of small cell sleep states are improving over time these gains could soon be much greater.

6.2 Future Work

While the contributions of this thesis were mainly centred on MU-MIMO-based reassignments of UEs between adjacent small cells, the methods used can be extended to other multi-antenna capability enhancements. In the following we identify and discuss a number of future works which will build upon the works of this thesis. The following list includes some works that we have already started while some correspond to works which can be pursued by other researchers.

6.2.1 Investigation of Additional Baselines for Ch5 Comparison

In Chapter 5 we compare the performance (in terms of energy and spectral efficiency) of our proposed MU-MIMO-based reassignment mechanism to the case in which UEs are simply assigned to the small cell eNB of highest RSRP. While this does represent standard practice [6] and, for our our scenario, correspond to the Signal to Interference and Noise Ratio (SINR)-maximising approach [7], we consider a potentially better baseline to be the case in which UEs can be reassigned to neighbouring small cells if their SINR increases as a result of the reassignment [115], given that their original cell is switched to a sleep state following the reassignment to conserve energy. In this case we expect that a combined approach in which it is possible to reassign based on this *or* our solution will provide the best overall performance.

6.2.2 Dynamic Investigation of Ch5's Long-term Performance

Chapter 5 investigates the energy savings that can be achieved by switching small cells to a sleep state as a result of them being emptied through MU-MIMO-based reassignments, starting from a situation in which all small cells are active. We do not, however, investigate the long-term case in which small cells will need to be reawoken and the configuration readjusted as a result of changing network dynamics. We expect the performance of this long-term case to be close to ours presented in this work although inefficiencies may also result due to the time taken to reawaken deactivated small cells.

6.2.3 Variation of Tx Power with Small Cell Density

In the scenarios investigated throughout this thesis we hold the small cell Transmit (Tx) powers constant as deployment densities (/Deployment Ratios (DRs)) increase. This results in higher inter-cell interference in more densely deployed scenarios. As a future work, for our proposed mechanisms, we could consider how system performance can be improved through enabling small cell power control in denser deployment scenarios. This is likely to result in a tradeoff between inter-cell interference and the amount of inter-cell coverage region overlap necessary for effective reassignments between adjacent small cells.

6.2.4 Coordinated Reassignments between MEA-equipped Small Cells

As stated earlier, the use of switched Multi-Element Antenna (MEA) systems is a beam-steering solution in which, multiple antennas, each of different antenna gain patterns (e.g. patch antennas – which essentially form a beam in a single 90° sector), point in different directions [5]. This provides a low-cost solution to small cell beam-steering as only a single transceiver chain is required. In these scenarios, however, in cases where no single antenna element (i.e. beam pattern) can serve all of the UEs assigned to a small cell at once, multiple antenna elements are activated, resulting in a loss in antenna gain and directionality.

For this reason we propose the use of a similar reassignment mechanism to the one proposed in Chapter 5, except instead of targeting the minimisation of the number of active small cells, we target the minimisation of the number of antenna elements required to serve the set of UEs assigned to each small cell. In this way we expect to remove a significant proportion of the losses associated with the simultaneous use of multiple antenna elements, while at the same time limiting interference between neighbouring cells.

6.2.5 Making Simulator Code Available

As identified in Chapter 3, in the course of this thesis a number of extensions to the Vienna System Level (SL) simulator [8] were required in order to enable the investigation of our proposed reassignment mechanism. As these extensions will be useful to others in future SL studies of multi-antenna capabilities in small cell deployment scenarios, they will be made available online.

6.2.6 Combined MU-MIMO Across Small Cells and CoMP Coordinated Scheduling

As brushed on in Chapter 4, as the Channel State Information (CSI) feedback required for *MU-MIMO across Small Cells* is a subset of that required for Coordinated Scheduling (CS-CoMP) the two methods could be performed in tandem, from which increased gains in spectral efficiency could be achieved.

Acronyms

3D 3 Dimensional

3GPP 3rd Generation Partnership Program

ABS Almost Blank Subframe

ARQ Automatic Repeat ReQuest

AWGN Additive White Gaussian Noise

BBU BaseBand Unit

B-CQI Best CQI

BD Block Diagonalisation

BICM Bit-Interleaved Coded Modulation

BLER BLock Error Rate

CBS-CoMP Coordinated Beam-Switching

CDI Channel Direction Indicator

CDF Cumulative Distribution Function

CEUSE Cell Edge User Spectral Efficiency

CLSM Closed-Loop Spatial Multiplexing

CO₂ Carbon Dioxide

CoMP Coordinated Multi-Point

- CQI** Channel Quality Indicator
- C-RAN** Cloud Radio Access Network
- CRS** Common Reference Signal
- CS-CoMP** Coordinated Scheduling
- CS/CB** Coordinated Scheduling/Coordinated Beamforming
- CSE** Cell Spectral Efficiency
- CSI** Channel State Information
- CSI-RS** CSI Reference Signal
- CSG** Closed Subscriber Group
- CS-RS** Cell Specific Reference Signal
- DCS** Dynamic Cell(/Point) Selection
- DM-RS** Demodulation Reference Signal
- DPC** Dirty Paper Coding
- DR** Deployment Ratio
- DSL** Digital Subscriber Line
- E³F** Energy Efficiency Evaluation Framework
- EARTH** Energy Aware Radio and neTwork tecHnologies
- EDN** Extremely Dense Network
- EESM** Exponential ESM
- eNB** evolved Node Base station
- EPC** Evolved Packet Core
- ESM** Effective SINR Mapping
- FDD** Frequency Division Duplexing

-
- FFR** Frequency Fractional Reuse
- FP7** European Commission 7th Framework Programme
- GUS** Greedy User Selection
- H-ARQ** Hybrid Automatic Repeat reQuest
- HeNB** Home/indoor femtocell eNB
- ICIC** Inter-Cell Interference Coordination
- IMT-A** International Mobile Telecommunications-Advanced
- InH** Indoor Hotspot
- IRC** Interference Rejection Combining
- JP** Joint Processing
- JT** Joint Transmission
- L2S** Link-to-System
- LESM** Logarithmic ESM
- LL** Link Level
- LOS** Line-Of-Sight
- LTE** Long Term Evolution
- LTE-A** LTE-Advanced
- MAC** Medium Access Control
- MCS** Modulation and Coding Scheme
- MEA** Multi-Element Antenna
- MESC** Maximum Expected SINR Combiner
- MIESM** Mutual Information Effective SINR Mapping
- MIESM** Mutual Information ESM

- MIMO** Multiple Input Multiple Output
- MME** Mobility Management Entity
- MMSE** Minimum Mean Square Error
- MNO** Mobile Network Operator
- MRC** Maximum Ratio Combining
- MU-MIMO** Multi-User MIMO
- MUI** Multi-User Interference
- NL** Network Level
- OLSM** Open-Loop Spatial Multiplexing
- OOP** Object Oriented Programing
- PDCP** Packet Data Convergence Protocol
- PF** Proportional Fair
- PHY** Physical
- PMI** Precoding Matrix Indicator
- PON** Passive Optical Network
- RB** Resource Block
- RI** Rank Indicator
- RE** Resource Element
- RF** Radio Frequency
- RLC** Radio Link Control
- RMa** Rural Macrocell
- ROI** Region of Interest
- RR** Round Robin

-
- RRC** Radio Resource Control
- RRH** Remote Radio Head
- RRM** Radio Resource Management
- RSRP** Reference Signal Received Power
- RUE** Reassigned UE
- Rx** Receive
- SAMURAI** Spectrum Aggregation and MU-MIMO: ReAl-world Impact
- S-GW** Serving GateWay
- SINR** Signal to Interference and Noise Ratio
- SISO** Single Input Single Output
- SL** System Level
- SIC** Successive Interference Cancellation
- SMa** Suburban Macrocell
- SNR** Signal to Noise Ratio
- SON** Self-Organising Networks
- SSPS** Semi-Static Point Selection
- SU-MIMO** Single-User MIMO beamforming on a single spatial stream
- SU-MIMO** Single-User MIMO
- SUS** Semi-orthogonal User Selection
- TM** Transmission Mode
- TTI** Transmission Time Interval
- TU** Typical Urban
- TUE** Target UE

Tx Transmit

TxD Transmit Diversity

UE User Equipment

ULA Uniform Linear Array

UMa Urban Macrocell

UMi Urban Microcell

URS UE-specific Reference Signal

WiSE Wireless System Engineering

ZF Zero-Forcing

ZFBF Zero-Forcing Beam Forming

Bibliography

- [1] D. Astely, E. Dahlman, G. Fodor, S. Parkvall, and J. Sachs, "LTE release 12 and beyond," *IEEE Communications Magazine*, vol. 51, no. 7, 2013.
- [2] D. Gesbert, M. Kountouris, R. Heath Jr., C.-b. Chae, and T. Salzer, "Shifting the MIMO Paradigm," *IEEE Signal Processing Magazine*, vol. 24, no. 5, pp. 36–46, 2007.
- [3] Small Cell Forum, "Small Cell Market Highlights 2014 Q1," Tech. Rep. June, 2014.
- [4] Small Cell Forum Press Release, "Small Cells Outnumber Traditional Mobile Base Stations," 2012. [Online]. Available: <http://www.smallcellforum.org/press-releases/small-cells-outnumber-traditional-mobile-base-stations/>
- [5] R. Razavi, L. Ho, H. Claussen, and D. Lopez-Perez, "Improving Small Cell Performance through Switched Multi-element Antenna Systems in Heterogeneous Networks," *IEEE Transactions on Vehicular Technology (TVT)*, vol. PP, no. 99, pp. 1–1, 2014.
- [6] 3GPP TS 36.304, "Evolved Universal Terrestrial Radio Access (E-UTRA); User Equipment (UE) procedures in idle mode," Tech. Rep.
- [7] A. Gupta, H. S. Dhillon, S. Vishwanath, and J. Andrews, "Downlink Multi-Antenna Heterogeneous Cellular Network with Load Balancing," *IEEE Transactions on Communications*, vol. 62, no. 11, pp. 1–1, 2014.
- [8] J. C. Ikuno, M. Wrulich, and M. Rupp, "System Level Simulation of LTE Networks," in *IEEE Vehicular Technology Conference (VTC-Spring)*, 2010.
- [9] Small Cell Forum, "Small Cell Definition." [Online]. Available: <http://www.smallcellforum.org/about/about-small-cells/small-cell-definition/>
- [10] S. Alamouti, "A simple transmit diversity technique for wireless communications," *IEEE Journal on Selected Areas in Communications (JSAC)*, vol. 16, no. 8, pp. 1451–1458, 1998.

- [11] A. Ghosh, J. Zhang, J. G. Andrews, and R. Muhammed, *Fundamentals of LTE*. Pearson Education, 2010.
- [12] L. Zheng and D. N. Tse, “Diversity and multiplexing: a fundamental tradeoff in multiple-antenna channels,” *IEEE Transactions on Information Theory*, vol. 49, no. 5, pp. 1073–1096, 2003.
- [13] 3GPP TR 25.943, “Technical Specification Group Radio Access Network; Deployment aspects,” Tech. Rep.
- [14] 3GPP TS 36.211, “Evolved Universal Terrestrial Radio Access (E-UTRA); Physical channels and modulation,” Tech. Rep.
- [15] 3GPP TS 36.213, “Evolved Universal Terrestrial Radio Access (E-UTRA); Physical layer procedures,” Tech. Rep.
- [16] C. Mehlführer *et al.*, “The Vienna LTE simulators - Enabling reproducibility in wireless communications research,” *EURASIP Journal on Advances in Signal Processing*, no. 1, p. 29, 2011.
- [17] J.-H. Noh and S.-J. Oh, “System-Level Simulation of LTE/LTE-A for IMT-Advanced Systems,” in *IEEE Vehicular Technology Conference (VTC-Spring)*, 2011, pp. 1–5.
- [18] S. Schwarz, M. Wrulich, and M. Rupp, “Mutual information based calculation of the Precoding Matrix Indicator for 3GPP UMTS/LTE,” in *IEEE International ITG Workshop on Smart Antennas (WSA)*, 2010, pp. 52–58.
- [19] A. Cattoni *et al.*, “Multi-user MIMO and Carrier Aggregation in 4G systems: The SAMURAI approach,” in *IEEE Wireless Communications and Networking Conference Workshops (WCNCW)*, Apr. 2012, pp. 288–293.
- [20] H. T. Nguyen and I. Z. Kovacs, “A MU-MIMO CQI Estimation Method for MU-MIMO UEs in LTE Systems,” *IEEE Vehicular Technology Conference (VTC-Fall)*, 2012.
- [21] M. H. Costa, “Writing on dirty paper (corresp.),” *IEEE Transactions on Information Theory*, vol. 29, no. 3, pp. 439–441, 1983.
- [22] G. Caire and S. Shamai, “On the achievable throughput of a multiantenna Gaussian broadcast channel,” *IEEE Transactions on Information Theory*, vol. 49, no. 7, pp. 1691–1706, 2003.
- [23] L.-U. Choi and R. Murch, “A Transmit Preprocessing Technique for Multiuser MIMO Systems Using a Decomposition Approach,” *IEEE Transactions on Wireless Communications (TWC)*, vol. 3, no. 1, pp. 20–24, 2004.

- [24] A. D. Dabbagh and D. J. Love, "Precoding for Multiple Antenna Gaussian Broadcast Channels With Successive Zero-Forcing," *IEEE Transactions on Signal Processing*, vol. 55, no. 7, pp. 3837–3850, 2007.
- [25] S. Schwarz and M. Rupp, "Antenna combiners for block-diagonalization based multi-user MIMO with limited feedback," in *IEEE International Conference on Communications (ICC) Workshops*, 2013, pp. 127–132.
- [26] G. Dimic and N. Sidiropoulos, "On downlink beamforming with greedy user selection: performance analysis and a simple new algorithm," *IEEE Transactions on Signal Processing*, vol. 53, no. 10, pp. 3857–3868, 2005.
- [27] T. Yoo, N. Jindal, and A. Goldsmith, "Multi-Antenna Downlink Channels with Limited Feedback and User Selection," *IEEE Journal on Selected Areas in Communications (JSAC)*, vol. 25, no. 7, pp. 1478–1491, 2007.
- [28] S. Sigdel, R. C. Elliott, W. A. Krzymien, and M. Al-Shalash, "Greedy and Genetic User Scheduling Algorithms for Multiuser MIMO Systems with Block Diagonalization," in *IEEE Vehicular Technology Conference (VTC-Fall)*, 2009, pp. 1–6.
- [29] N. Ravindran and N. Jindal, "Multi-user Diversity vs. Accurate Channel State Information - MATLAB Code," 2009. [Online]. Available: www.ece.umn.edu/users/nihar/mud_csi_code.html
- [30] N. Ravindran and N. Jindal, "Multi-User Diversity vs. Accurate Channel State Information in MIMO Downlink Channels," *IEEE Transactions on Wireless Communications (TWC)*, vol. 11, no. 9, pp. 3037–3046, 2012.
- [31] T. Minka, "Lightspeed Matlab Toolbox v2.6." [Online]. Available: <http://research.microsoft.com/en-us/um/people/minka/software/lightspeed/>
- [32] D. P. Murphy, D. Finn, N. Marchetti, and L. A. DaSilva, "Performance-Complexity Tradeoff in Multi-User MIMO Scheduling," Tech. Rep., 2014.
- [33] S. Venkatesan, A. Lozano, and R. Valenzuela, "Network MIMO: Overcoming Intercell Interference in Indoor Wireless Systems," in *Asilomar Conference on Signals, Systems and Computers (ACSSC)*, 2007, pp. 83–87.
- [34] China Mobile Research Institute, "C-RAN The Road Towards Green RAN," Beijing, China, Tech. Rep., 2011.
- [35] Nokia Siemens Networks, "Liquid Radio: Let traffic waves flow most efficiently," 2013. [Online]. Available: http://nsn.com/system/files/document/nokia_siemens_networks_liquid_radio_white_paper_05_06_2013_0.pdf

- [36] 3G Americas, “3GPP Mobile Broadband Innovation Path to 4G: Release 9, Release 10 and Beyond,” 2010. [Online]. Available: www.4gamericas.org/documents/3GPP_Rel-9_BeyondFeb2010.pdf
- [37] 3GPP TR 36.819, “Coordinated multi-point operation for LTE physical layer aspects,” Tech. Rep.
- [38] P. Marsch and G. P. Fettweis, *Coordinated Multi-Point in Mobile Communications From Theory to Practice*. Cambridge University Press, 2011.
- [39] M. Dohler, R. W. Heath, A. Lozano, C. B. Papadias, and R. A. Valenzuela, “Is the PHY layer dead?” *IEEE Communications Magazine*, vol. 49, no. 4, pp. 159–165, 2011.
- [40] E. Pateromichelakis, M. Shariat, A. ul Quddus, and R. Tafazolli, “On the Evolution of Multi-Cell Scheduling in 3GPP LTE / LTE-A,” *IEEE Communications Surveys & Tutorials*, vol. 15, no. 2, pp. 701–717, 2013.
- [41] M. Sawahashi, Y. Kishiyama, A. Morimoto, D. Nishikawa, and M. Tanno, “Coordinated multipoint transmission/reception techniques for LTE-advanced [Coordinated and Distributed MIMO],” *IEEE Wireless Communications*, vol. 17, no. 3, pp. 26–34, 2010.
- [42] P. Hosein and C. van Rensburg, “On the Performance of Downlink Beamforming with Synchronized Beam Cycles,” in *IEEE Vehicular Technology Conference (VTC-Spring)*, 2009.
- [43] C. Van Rensburg and P. Hosein, “Interference Coordination through Network-Synchronized Cyclic Beamforming,” in *IEEE Vehicular Technology Conference (VTC-Fall)*, 2009, pp. 1–5.
- [44] J. Ellenbeck, M. Hammoud, B. Lazarov, and C. Hartmann, “Autonomous beam coordination for the downlink of an IMT-Advanced cellular system,” in *IEEE European Wireless Conference (EW)*, 2010, pp. 602–607.
- [45] A. Gamst and K. Ralf, “Computational complexity of some interference graph calculations (mobile radio),” *IEEE Transactions on Vehicular Technology (TVT)*, vol. 39, no. 2, pp. 140–149, 1990.
- [46] L. Liu, J. C. Zhang, J.-C. Yu, and J. Lee, “Intercell Interference Coordination through Limited Feedback,” *International Journal of Digital Multimedia Broadcasting*, 2010.
- [47] M. Feng, X. She, L. Chen, and Y. Kishiyama, “Enhanced Dynamic Cell Selection with Muting Scheme for DL CoMP in LTE-A,” in *IEEE Vehicular Technology Conference (VTC-Spring)*, 2010.

- [48] K. Lee and D. Williams, "A space-frequency transmitter diversity technique for OFDM systems," in *IEEE Global Telecommunications Conference (Globecom)*, vol. 3, 2000, pp. 1473–1477.
- [49] H. Huang and S. Venkatesan, "Asymptotic downlink capacity of coordinated cellular networks," in *IEEE Asilomar Conference on Signals, Systems and Computers (ACSSC)*, vol. 1, 2004, pp. 850–855.
- [50] D. Gesbert *et al.*, "Multi-Cell MIMO Cooperative Networks: A New Look at Interference," *IEEE Journal on Selected Areas in Communications (JSAC)*, vol. 28, no. 9, pp. 1380–1408, 2010.
- [51] S.-N. Hong and G. Caire, "Compute-and-Forward Strategies for Cooperative Distributed Antenna Systems," *IEEE Transactions on Information Theory*, vol. 59, no. 9, pp. 5227–5243, 2013.
- [52] Y.-H. Nam *et al.*, "Evolution of reference signals for LTE-advanced systems," *IEEE Communications Magazine*, vol. 50, no. 2, pp. 132–138, 2012.
- [53] H. Holma and A. Toskala, *LTE-Advanced: 3GPP Solution for IMT-Advanced*. Wiley, 2012.
- [54] S. Ye, S. H. Wong, and C. Worrall, "Enhanced physical downlink control channel in LTE advanced Release 11," *IEEE Communications Magazine*, vol. 51, no. 2, pp. 82–89, 2013.
- [55] H. Claussen, L. T. W. Ho, and L. G. Samuel, "An overview of the femtocell concept," *Bell Labs Technical Journal*, vol. 13, no. 1, pp. 221–245, 2008.
- [56] Small Cell Forum, "SCF100.04.03: Release Structure and Roadmap," Tech. Rep. June, 2014.
- [57] T. Q. Quek, G. de la Roche, I. Güvenç, and M. Kountouris, *Small cell networks: Deployment, PHY techniques, and resource management*. Cambridge University Press, 2013.
- [58] H. Claussen and F. Pivit, "Femtocell Coverage Optimization Using Switched Multi-Element Antennas," in *IEEE International Conference on Communications (ICC)*, 2009, pp. 1–6.
- [59] R. Razavi and H. Claussen, "Self-configuring Switched Multi-Element Antenna system for interference mitigation in femtocell networks," in *IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, 2011, pp. 237–242.

- [60] D.-C. Oh, H.-C. Lee, and Y.-H. Lee, "Power Control and Beamforming for Femtocells in the Presence of Channel Uncertainty," *IEEE Transactions on Vehicular Technology (TVT)*, vol. 60, no. 6, pp. 2545–2554, 2011.
- [61] V. Chandrasekhar, M. Kountouris, and J. Andrews, "Coverage in multi-antenna two-tier networks," *IEEE Transactions on Wireless Communications (TWC)*, vol. 8, no. 10, pp. 5314–5327, 2009.
- [62] D. Jaramillo-Ramirez, M. Kountouris, and E. Hardouin, "Downlink beamforming in multi-antenna two-tier networks with user selection," in *IEEE GLOBECOM Workshops (GC Wkshps)*, 2011, pp. 258–263.
- [63] A. A. Dowhuszko, F. Ahmed, and O. Tirkkonen, "Decentralized Transmit Beamforming scheme for interference coordination in small cell networks," in *IEEE International Black Sea Conference on Communications and Networking (BlackSeaCom)*, 2013, pp. 121–126.
- [64] JDSU - White Paper, "Cloud-RAN Deployment with CPRI Fronthaul Technology," Tech. Rep., 2013. [Online]. Available: http://www.jdsu.com/ProductLiterature/cloudRAN_wp_tfs_nse_ae.pdf
- [65] S.-H. Park, O. Simeone, O. Sahin, and S. S. (Shitz), "Robust and Efficient Distributed Compression for Cloud Radio Access Networks," *IEEE Transactions on Vehicular Technology (TVT)*, vol. 62, no. 2, pp. 692–703, 2013.
- [66] Alcatel Lucent, "lightRadio™: Evolve your wireless broadband network." [Online]. Available: <http://www.alcatel-lucent.com/solutions/lightradio>
- [67] J. Beas, G. Castanon, I. Aldaya, A. Aragon-Zavala, and G. Campuzano, "Millimeter-Wave Frequency Radio over Fiber Systems: A Survey," *IEEE Communications Surveys & Tutorials*, vol. 15, no. 4, pp. 1593–1619, 2013.
- [68] P. Gamage, A. Nirmalathas, C. Lim, D. Novak, and R. Waterhouse, "Design and Analysis of Digitized RF-Over-Fiber Links," *IEEE/OSA Journal of Lightwave Technology (JLT)*, vol. 27, no. 12, pp. 2052–2061, 2009.
- [69] "Ericsson Radio Dot System," 2013. [Online]. Available: <http://www.ericsson.com/ourportfolio/products/radio-dot-system>
- [70] T. Nakamura *et al.*, "Trends in small cell enhancements in LTE advanced," *IEEE Communications Magazine*, vol. 51, no. 2, pp. 98–105, 2013.
- [71] V. Chandrasekhar, J. Andrews, and A. Gatherer, "Femtocell networks: a survey," *IEEE Communications Magazine*, vol. 46, no. 9, pp. 59–67, 2008.

- [72] M. Z. Shakir *et al.*, “Green heterogeneous small-cell networks: toward reducing the CO₂ emissions of mobile communications industry using uplink power adaptation,” *IEEE Communications Magazine*, vol. 51, no. 6, pp. 52–61, 2013.
- [73] “FP7 Energy Aware Radio and neTwork tecHnologies (EARTH) Project.” [Online]. Available: www.ict-earth.eu
- [74] M. A. Imran *et al.*, “FP7 EARTH D3.2: Energy efficiency analysis of the reference systems , areas of improvements and target breakdown,” Tech. Rep., 2010. [Online]. Available: http://ec.europa.eu/information_society/apps/projects/logos/3/247733/080/deliverables/001_EARTHWP2D23v2.pdf
- [75] G. Auer *et al.*, “How much energy is needed to run a wireless network?” *IEEE Wireless Communications*, vol. 18, no. 5, pp. 40–49, 2011.
- [76] S. Auroux, M. Dräxler, A. Morelli, and V. Mancuso, “CROWD Dynamic Network Reconfiguration in Wireless DenseNets,” in *European Conference on Networks and Communications (EuCNC)*, 2014.
- [77] “FP7 METIS Project.” [Online]. Available: www.metis2020.com
- [78] G. Piro, N. Baldo, and M. Miozzo, “An LTE module for the ns-3 network simulator,” in *International ICST Conference on Simulation Tools and Techniques (SIMUTools)*, 2011, pp. 415–422.
- [79] K. Brueninghaus *et al.*, “Link Performance Models for System Level Simulations of Broadband Radio Access Systems,” in *IEEE International Conference on Personal, Indoor, and Mobile Radio Communications (PIMRC)*, vol. 4, 2005, pp. 2306–2311.
- [80] P. Mogensen *et al.*, “LTE Capacity Compared to the Shannon Bound,” in *IEEE Vehicular Technology Conference (VTC-Spring)*, 2007, pp. 1234–1238.
- [81] V. Chandrasekhar and J. G. Andrews, “Spectrum allocation in tiered cellular networks,” *IEEE Transaction on Communications*, vol. 57, no. 10, pp. 3059 – 3068, 2009.
- [82] “Google summer of code - accepted projects.” [Online]. Available: <http://www.nsnam.org/wiki/GSOC2014AcceptedProjects>
- [83] J. Andrews, “Seven ways that hetnets are a cellular paradigm shift,” *IEEE Communications Magazine*, vol. 51, no. 3, pp. 136–144, 2013.
- [84] X. Wu *et al.*, “An XG-PON module for the ns-3 network simulator,” in *International ICST Conference on Simulation Tools and Techniques (SIMUTools)*, 2013, pp. 195–202.

- [85] R. Gupta *et al.*, “Labview based platform for prototyping dense lte networks in crowd project.” [Online]. Available: http://www.ict-crowd.eu/downloads/2014/LabVIEW_Platform_prototyping_LTE.pdf
- [86] A. Alexandrescu, *Modern C++ Design: Generic Programming and Design Patterns Applied*. Addison Wesley, 2001.
- [87] 3GPP TS 36.814, “Further Advancements for E-UTRA Physical Layer Aspects,” Tech. Rep.
- [88] M. Trivellato, F. Boccardi, and H. Huang, “On transceiver design and channel quantization for downlink multiuser MIMO systems with limited feedback,” *IEEE Journal on Selected Areas in Communications (JSAC)*, vol. 26, no. 8, pp. 1494–1504, 2008.
- [89] R1-093740, “Overhead Assumption for Performance Evaluation in ITU-R Submission,” NTT DOCOMO, Tech. Rep.
- [90] Z. Sun, C. Yin, and G. Yue, “Reduced-Complexity Proportional Fair Scheduling for OFDMA Systems,” in *IEEE International Conference on Communications, Circuits and Systems (ICCCAS)*, 2006.
- [91] L. Liu *et al.*, “Downlink MIMO in LTE-advanced: SU-MIMO vs. MU-MIMO,” *IEEE Communications Magazine*, vol. 50, no. 2, pp. 140–147, 2012.
- [92] S. Schwarz, C. Mehlhruer, and M. Rupp, “Low complexity approximate maximum throughput scheduling for LTE,” in *IEEE Signals, Systems and Computers (ASILOMAR)*, 2010, pp. 1563–1569.
- [93] L. F. Del Carpio Vega, “System level modeling and evaluation of advanced linear interference aware receivers,” M.Sc., School of Electrical Engineering, Aalto University, Finland, 2012.
- [94] ITU-R M.2135, “Guidelines for evaluation of radio interface technologies for IMT-Advanced,” Tech. Rep., 2009.
- [95] 3GPP RAN 4, “Simulation assumptions and parameters for FDD HeNB RF requirements.” R4-092042, May 2009.
- [96] METIS, “Deliverable D6 . 1 Simulation guidelines,” Tech. Rep., 2013.
- [97] Ericsson, “Radio Dot System Telebriefings,” 2013. [Online]. Available: www.ericsson.com/res/investors/docs/2013/ericsson-radio-dot-system-telebriefings-26-sept.pdf
- [98] P. D. Francesco, F. Malandrino, and L. A. DaSilva, “Mobile Network Sharing Between Operators: A Demand Trace-Driven Study,” in *ACM SIGCOMM Capacity Sharing Workshop (CSWS)*, 2014.

- [99] J. Kibilda and L. A. DaSilva, "Efficient coverage through inter-operator infrastructure sharing in mobile networks," in *IEEE IFIP Wireless Days (WD)*, 2013, pp. 1–6.
- [100] M. Wellens, J. Riihijärvi, and P. Mähönen, "Empirical time and frequency domain models of spectrum use," *Physical Communication*, vol. 2, no. 1-2, pp. 10–32, 2009.
- [101] I. Macaluso, D. Finn, B. Ozgul, and L. A. DaSilva, "Complexity of Spectrum Activity and Benefits of Reinforcement Learning for Dynamic Channel Selection," *IEEE Journal on Selected Areas in Communications (JSAC)*, vol. 31, no. 11, pp. 2237–2248, 2013.
- [102] S. Fortune *et al.*, "WISE design of indoor wireless systems: practical computation and optimization," *IEEE Computational Science and Engineering*, vol. 2, no. 1, pp. 58–68, 1995.
- [103] D. Finn, H. Ahmadi, A. Cattoni, and L. A. DaSilva, "Improved spectral efficiency through multi-user mimo across small cells," 2015, submitted to *IEEE Transactions on Vehicular Technology (TVT)*.
- [104] D. Finn, H. Ahmadi, A. Cattoni, and L. A. DaSilva, "Multi-User MIMO across Small Cells," in *IEEE International Conference on Communications (ICC)*, 2014.
- [105] "Small Cell Forum Market Status Reports." [Online]. Available: <http://www.smallcellforum.org/resources-reports>
- [106] J. Duplicy *et al.*, "MU-MIMO in LTE Systems," *EURASIP Journal on Wireless Communications and Networking (JWCN)*, vol. 2011, no. 1, pp. 1–13.
- [107] H. Huang and M. Trivellato, "Performance of multiuser MIMO and network coordination in downlink cellular networks," in *IEEE International Symposium on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt)*, 2008, pp. 85–90.
- [108] P. Kyösti *et al.*, "Matlab Implementation of the WINNER Phase II Channel Model ver1.1," Tech. Rep., 2007. [Online]. Available: http://www.ist-winner.org/phase_2_model.html
- [109] 3GPP TS 36.133, "Evolved Universal Terrestrial Radio Access (E-UTRA); Requirements for support of radio resource management," Tech. Rep.
- [110] 3GPP TS 36.331, "Evolved Universal Terrestrial Radio Access (E-UTRA); Radio Resource Control (RRC); Protocol specification," Tech. Rep.

- [111] D. Finn, H. Ahmadi, R. Razavi, H. Claussen, and L. A. DaSilva, "Energy and spectral efficiency gains from multi-user mimo-based small cell reassignments," 2015, submitted to *IEEE Global Communications Conference (Globecom)*.
- [112] I. Ashraf, F. Boccardi, and L. Ho, "SLEEP mode techniques for small cell deployments," *IEEE Communications Magazine*, vol. 49, no. 8, pp. 72–79, 2011.
- [113] G. Biczók, J. Malmodin, and A. Fehske, "FP7 EARTH Deliverable 2.1: Economic and Ecological Impact of ICT," Tech. Rep., 2011. [Online]. Available: https://bscw.ict-earth.eu/pub/bscw.cgi/d38532/EARTH_WP2_D2.1_v2.pdf
- [114] I. Ashraf, L. T. W. Ho, and H. Claussen, "Improving Energy Efficiency of Femtocell Base Stations Via User Activity Detection," in *IEEE Wireless Communications and Networking Conference (WCNC)*, 2010, pp. 1–5.
- [115] Y. Li, Y. Jia, Y. Wang, and Q. Liu, "Collaborative Sleeping Scheme for Femtocell Networks," in *IEEE International Conference on Green Computing and Communications (GreenCom), Internet of Things (iThings) and Cyber, Physical and Social Computing (CPSCom)*, 2013, pp. 142–147.
- [116] C. Bottai, C. Cicconetti, A. Morelli, M. Rosellini, and C. Vitale, "Energy-Efficient User Association In Extremely Dense Small Cells," in *European Conference on Networks and Communications (EuCNC)*, 2014.
- [117] S. Jiang, F. R. Yu, and Y. Sun, "Distributed energy-efficient inter-cell interference control with BS sleep mode and user fairness in cellular networks," in *2013 IEEE Global Communications Conference (GLOBECOM)*. IEEE, Dec. 2013, pp. 2581–2586.
- [118] M. Z. Shakir, D. Lopez-Perez, M. A. Imran, and K. A. Qaraqe, "Tutorial #5: Small-cells: Capacity, Mobility and Energy Efficiency Perspectives," in *IEEE International Conference on Communications (ICC)*, Sydney, Australia, 2014.
- [119] V. Chvatal, "A greedy heuristic for the set-covering problem," *Mathematics of operations research*, vol. 4, no. 3, pp. 233–235, 1979.
- [120] D. Prajapati and V. Richhariya, "A Survey on Cell Selection Schemes for Femtocell Networks," vol. 5013, no. 3, pp. 465–468, 2014.
- [121] A. Mesodiakaki, F. Adelantado, L. Alonso, and C. Verikoukis, "Energy-efficient user association in cognitive heterogeneous networks," *IEEE Communications Magazine*, vol. 52, no. 7, pp. 22–29, 2014.
- [122] L. S. Mohjazi, M. a. Al-Qutayri, H. R. Barada, K. F. Poon, and R. M. Shubair, "Self-optimization of pilot power in enterprise femtocells using multi objective heuristic," *Journal of Computer Networks and Communications*, vol. 2012, 2012.

-
- [123] F. Mhiri, K. Sethom Ben Reguiga, R. Bouallegue, and G. Pujolle, "A power management algorithm for green femtocell networks," in *2011 The 10th IFIP Annual Mediterranean Ad Hoc Networking Workshop*. IEEE, Jun. 2011, pp. 45–49.
- [124] H. Kim and T. Chen, "Energy efficient femtocell power management," in *2012 International Symposium on Wireless Communication Systems (ISWCS)*. IEEE, Aug. 2012, pp. 586–590.