Scheduling Strategies and Interference Mitigation for OFDMA Cellular Networks

submitted by

Ashley Mills

for the degree of Doctor of Engineering

of the

University of Bath

Department of Computer Science

September 2011

Signature of Author ....

Ashley Mills

A few words

Some people claim to be “Self Made” but the truth is nobody is.

Each and every one depends on each and every other.

To the people that helped and supported me through this EngD.

Truly, thank you all.

Summary

The next generation of cellular networks will employ OFDMA (Orthogonal Frequency Division Multiple Access) to deliver high rate data services to a range of users expressing heterogeneous demands. This dissertation looks at MAC (Medium Access Control) layer techniques for improving the experience of these users, with the goal of decoupling their experience from their geographical positions. The work is undertaken with a view to minimising the impact of suggested solutions on peak cell throughputs and other measures of service quality. The primary agents by which the goal is achieved are: frequency selective scheduling and interference mitigation.

The performance of frequency selective scheduling is first analysed under 802.16e Mobile WiMAX, for different user velocities, with the goal of understanding its relevance in extant deployments. It is shown to exceed the efficiency of alternatives over the range of velocities experienced by typical users, leading to the conclusion that it should be used by default in the majority of cells. The work furthers understanding of the relationship between user velocity and frequency selectivity under realistic conditions.

The dissertation progresses to enhancing the gains of frequency selective scheduling through augmentation with interference mitigation. The investigation proceeds by a comprehensive analysis of a static interference mitigation approach called soft-reuse, and identifies its determinant performance factors. The analysis demonstrates that static schemes are inappropriate for the types of dynamism seen in real networks, and motivates the need to develop schemes that can adapt to network changes.

Finally, the dissertation reviews dynamic approaches to interference avoidance and derives a novel intrasite interference mitigation scheme, that leverages the knowledge obtained earlier from the work on frequency selective scheduling and static interference mitigation. The scheme obtains robust albeit modest performance gains against a contemporary sector-independent frequency selective scheduling approach. The gains are evident whenever the cells at a site are not fully loaded, and when scheduling is not overly biased to users with low SINR. The work recommends that deployments adopt intrasite scheduling.

Contents

List of Acronyms

2D	Two Dimensional
3D	Three Dimensional
3G	3rd Generation mobile telecommunications
3GPP	3rd Generation Partnership Project
3GPP2	3rd Generation Partnership Project Two
AMC	Adaptive Modulation and Coding
ARQ	Automatic Repeat reQuest
AWGN	Additive White Gaussian Noise
BS	Base Station
BPS	Bits Per Symbol
BPSK	Binary Phase Shift Keying
BW	Bandwidth
CAPEX	CAPital EXpenditure
CDF	Cumulative Distribution Function
CDMA	Code Division Multiple Access
CQI	Channel Quality Index
CINR	Carrier to Interference + Noise Ratio
CoMP	Coordinated Multi Point Transmission / Reception
CO2	Carbon Dioxide
DC	Direct Current
DL	Downlink
DL-MAP	Downlink MAP
GPRS	General Packet Radio Service
GSM	Global System for Mobile Communications, originally Groupe Spécial Mobile
EDGE	Enhanced Data rates for GSM Evolution
EPSRC	Engineering and Physical Sciences Research Council
FCH	Frame Control Header
FFT	Fast Fourier Transfer
FDD	Frequency Division Duplex
FDM	Frequency Division Multiplexing
FDMA	Frequency Division Multiple Access
FEC	Forward Error Correction

FSS	Frequency Selective Scheduling
FUSC	Full Use of Sub Carriers
eNB	Evolved Node B
HARQ	Hybrid Automatic Repeat reQuest
HII	High Interference Indicator
HSPA	High Speed Packet Access
HSDPA	High Speed Downlink Packet Access
HSUPA	High Speed Uplink Packet Access
ICIC	Inter-Cell Interference Coordination
ICO	Interference COordination
IEEE	Institute of Electrical and Electronics Engineers
IMT	International Mobile Telecommunications
IP	Internet Protocol
IPR	Intellectual PRoperty
ITU	International Telecommunication Union
LTE	Long Term Evolution
MAC	Medium Access Control
MCS	Modulation and Coding Scheme
MIMO	Multiple Input Multiple Output
MONOTAS	Mobile Network Optimisation Through Advanced Simulation
MS	Mobile Station
MSE	Mean Squared Error
NIMBY	Not In My Back Yard
NME	Network Management Entity
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency-Division Multiple Access
OI	Overload Indicator
OPEX	OPerational EXpenditure
PC	Physical Cluster
PDU	Protocol Data Unit
PHY	Physical Layer
PRB	Physical Resource Block
PUSC	Partial Use of Sub Carriers
QoS	Quality of Service

QAM	Quadrature Amplitude Modulation
QFD	Quality Function Deployment
QPSK	Quadrature Phase Shift Keying
RC	Renumbered Cluster
R&D	Research and Development
RE	Research Engineer
RF	Radio Frequency
RNC	Radio Network Controller
RNTP	Relative Narrowband Transmit Power
RSRP	Reference Signal Received Power
RSSI	Received Signal Strength Indicator
RTG	Receive/transmit Transition Gap
SCM	Spatial Channel Model
SCME	Spatial Channel Model (Extended)
SIM	Subscriber Identity Module
SINR	Signal to Interference plus Noise Ratio
SNR	Signal to Noise Ratio
SM	Spatial Multiplexing
SOCRATES	Self-Optimisation and self-ConfiguRATion in wirelEss networkS
SON	Self Organising Network functionality
STBC	Space-Time Block Code
TDD	Time Division Duplex
TDM	Time Division Multiplexing
TDMA	Time Division Multiple Access
TTI	Transmission Time Interval
TTG	Transmit/receive Transition Gap
UE	User Equipment
UL	Uplink
UL-MAP	Uplink MAP
UMTS	Universal Mobile Telecommunications System
UTRA	UMTS Terrestrial Radio Access
UTRAN	UMTS Terrestrial Radio Access Network
VOIP	Voice Over IP
VRB	Virtual Resource Block

WiMAX	Worldwide Interoperability for Microwave Access

Chapter 1

Introduction

"*The future belongs to mobile internet: Vodafone is investing in capacity and quality Internet for all soon to be a reality.*" - Vodafone Press Release [1].

The mobile phone industry is currently, and in the absence of abating demand is perpetually, focused on building the next generation of communications technology. Internally within Vodafone (the company), the current target of this process has been given the name "The Network of the Future". The future in this context typically means the next 5 to 10 years. The principle drivers for this programme are the increased demand for novel high bandwidth mobile services coupled with the impending exhaustion of existing capacity.

Relative to the extant technology, the future network promises to offer: higher peak throughputs, lower latencies, higher spectral efficiencies, higher energy efficiencies, reduced complexity, and greater resilience to change.

The latter two goals subsume desires to automate many currently manual configuration and maintenance tasks, as well as to adapt dynamically and automatically in the face of ephemeral network artifacts that impact performance. Satisfaction of these goals is expected to confer the benefits of reduced OPEX (Operating Expenditure) and improved quality of service.

The task is complicated by a future network, envisaged to comprise of a multitude and manifold increase in, heterogeneous cells, services, and user capabilities spanning disparate and overlapping spatial extents from the home through to the macro.

This dissertation is concerned with homogenising user experience in the future network, in the face of expected diverse user geography, radio conditions, and demand. It achieves this through practical techniques that reduce the impact of inter-cell interference, united with scheduling strategies that exploit diversity that would otherwise act as a hindrance.

The rest of this chapter is organised as follows: in the next two Sections (1.1 and 1.2) a holistic overview of the problem domain is sketched in relation to Vodafone's operational objectives and wider organisational constraints and imperatives. The intention is to frame the focused research work that follows in the context of the "complete system" that is Vodafone, such that the multitude of frequently inimical stakeholder objectives, that often have a bearing on technical decisions, may be properly understood.

Following this, in Section 1.3, for each of the three main research contributions of this dissertation that appear in Chapters 3, 4, and 5, respectively, the narrower embodied industrial context of each piece of work is introduced and summarised in chronological order, described with respect to other events and trends within the ongoing operation of Vodafone R&D.

Finally, Section 1.4 summarises the contributions of this chapter, and outlines the structure of the dissertation remainder.

1.1 Understanding the Whole Problem

The systems engineer must understand the whole problem before trying to solve it [2] and is required to face in multiple directions [3]. This latter point is described clearly in [4] in the sentence "Systems engineering seeks a safe and balanced design in the face of opposing interests and multiple, sometimes conflicting constraints". This section examines the whole problem and conflicting constraints faced by the mobile communications engineer. It provides the systems context for the research work which follows.

The whole problem can be partitioned into two major tiers. The lower tier consists of the core network itself and is primarily concerned with physical elements of the network. The upper tier contains the operating companies, organisational bodies, and management interfaces. Figure 1-1 presents the lower tier in the form of a Checkland inspired rich picture [5, 6].

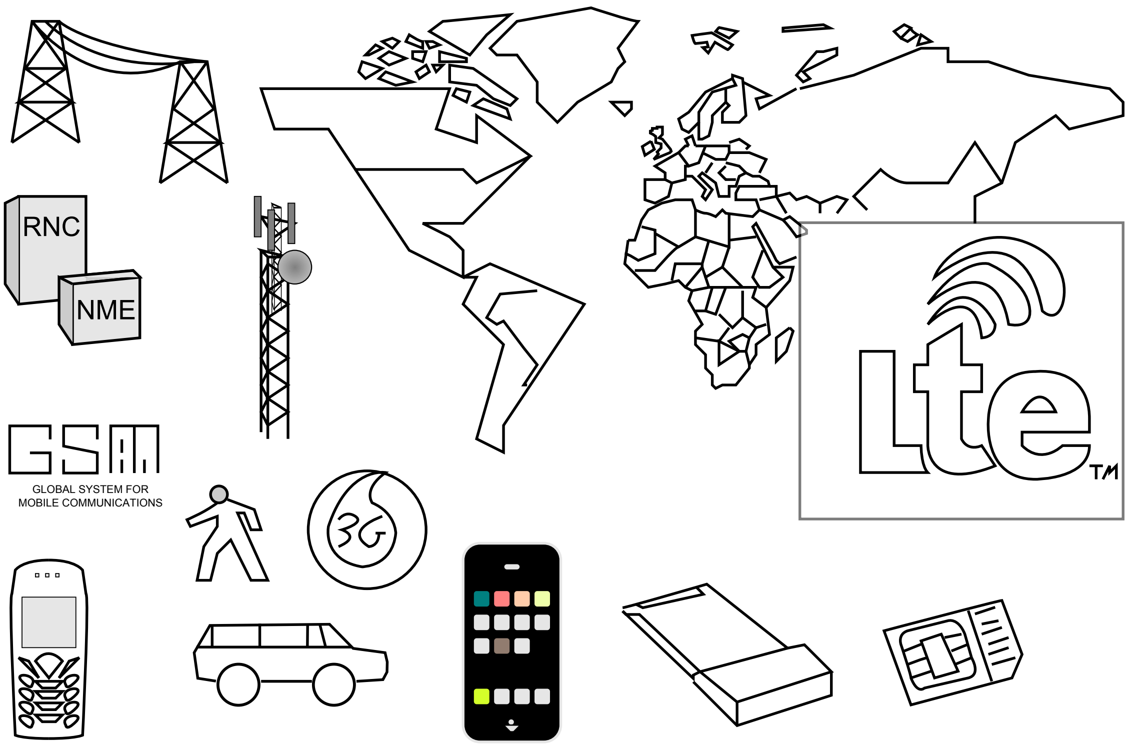

Figure 1-1: The technical aspects of the "network of the future", a rich picture.

The figure illustrates a number of important concepts regarding the core network. A prominent component of the picture is the world map. This is a reminder that the Vodafone network is a global network, and as such, global interoperability is a concern. At a technical level, handsets have to support at least ten operating bands, which require power hungry RF (Radio Frequency) chips and manifest costly patent fees. At the organisational level, each country has its own economic structure and banking laws, regulatory bodies and civil laws, as well as consumer fashions and cultural differences. These must all be observed, and the impact on technical aspects examined.

The LTE (Long Term Evolution) [1] logo features prominently to the right of the picture, because it is the dominant future technology currently being pursued by Vodafone as well

[1] LTE is the successor to the technology given the vernacular name "3G"

as the object of study in the majority of this dissertation.

To the left of the world map, various pieces of network equipment can be seen. The power pylons remind us that mobile networks consume huge amounts of power, a logistics problem in itself, and forever a concern, especially as the impact of natural resource shortages and global warming take their toll on business profits and generate new environmental legislature. A transmitting tower is shown to represent the tens of thousands of basestations that the company operates as a seamless whole. The RNC (Radio Network Controller) and NME (Network Management Entity) remind us that much of the hardware running the system, whilst hidden away in operating rooms, nevertheless must interact with the rest of the network seamlessly and withstand evolving network demands.

Below this, symbols for GSM (Global System for Mobile Communications, originally Groupe Spécial Mobile) [7, 8, 9] and 3G (3rd Generation mobile communications) [10, 11, 12] illustrate that the network concurrently supports these two network standards, and will continue to do so for the foreseeable future: legacy equipment is always an issue. New technology must not damage the operation of existing technology, and may be restricted in its development and deployment accordingly.

Along the bottom row, two different handsets are shown to illustrate the old yet persistent voice and text messaging centric GRPS (General Packet Radio Service) and EDGE (Enhanced Data rates for GSM Evolution) services, and the new data oriented 3G / HSPA (High Speed Packet Access) services, the laptop data card shown emphasises the importance and continued growth of data traffic. Users have different capabilities and demands. The pedestrian and car symbols along with the datacard, represent different levels of user mobility that the system must simultaneously support. Finally a SIM (Subscriber Identity Module) card is shown on the far right to signify privacy and security. The whole Vodafone network uses encrypted communication, and reminds us that security and fraud are always prominent concerns.

Figure 1-2 illustrates the outer layer of network interactions in the form of a rich picture.

The central cloud shows the previous rich picture (the network of the future) next to the Vodafone logo. The arrows between the logo and the network of the future, indicate the work that Vodafone is engaged in engineering the latter as a focus of Group R&D activities. Many aspects impinge upon developing this goal, as indicated by the surrounding items.

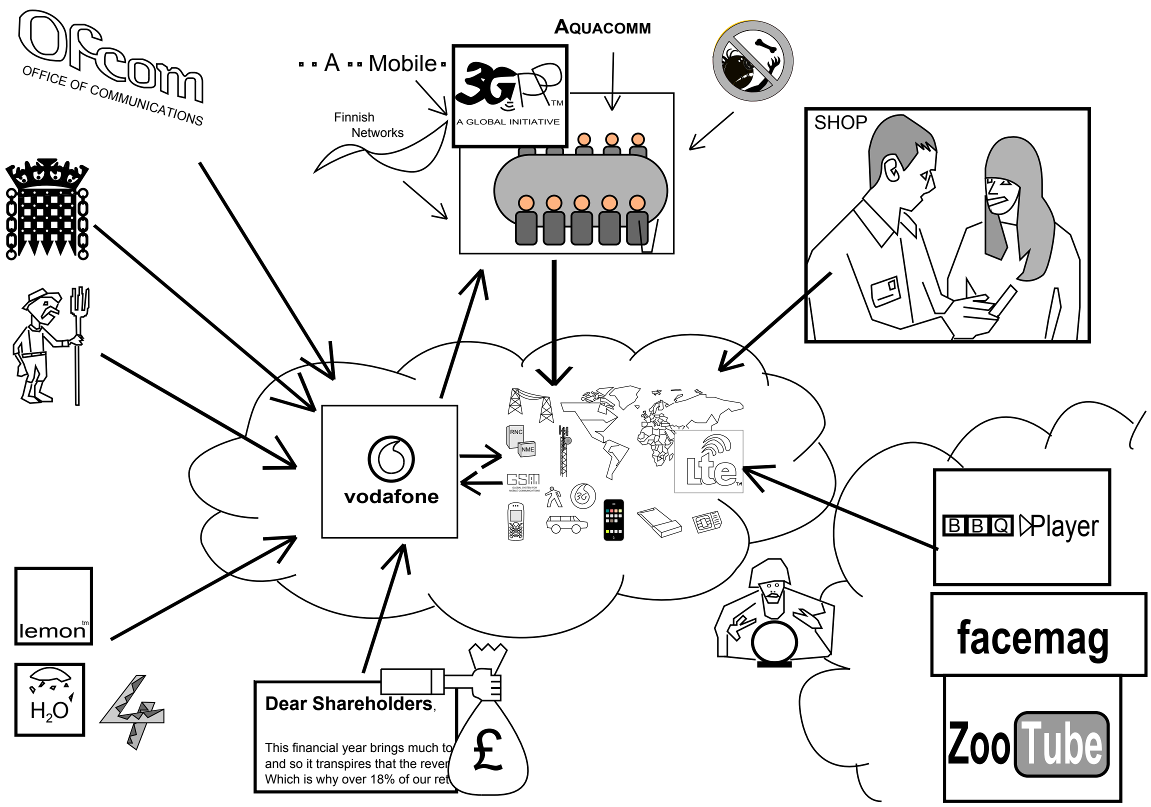

Figure 1-2: The organisational aspects of the “network of the future”, a rich picture.

The central top cluster represents the standards bodies. LTE is being developed by a global standards body called the 3GPP (3rd Generation Partnership Project) [13]. 3GPP has many partners, who are all trying to get their design decisions to be standardised. This includes hardware vendors, software vendors, other operators, patent trolls, and other groups. The standards impose restrictions on what is technically possible in LTE, and thus must be monitored closely. Furthermore, if changes are needed to the standards to improve support for a customer requirement, then it is the duty of the operator to try and push those changes through.

At the top left various external organisations are shown. Ofcom [14] is the independent regulatory body for the wireless communications industry. Ofcom are involved in many regulatory aspects, but most importantly to Vodafone they are responsible for licensing radio spectrum, holding spectrum auctions, and imposing licensing conditions [15]. The spectrum that Vodafone purchases has an impact on design decisions, since differences in physical behaviour of radiowaves at different frequencies is significant. Furthermore, the amount of bandwidth licensed restricts the type of radio bearer and consequent service that can be offered. Thus the operator must make the right decisions when it comes to purchasing spectrum, based on the expected usage, and must pay the right price to ensure a profitable future.

Below Ofcom is the symbol for the government, to indicate that Vodafone must play by the rules imposed by the country in which it is operating. Below the government is an "angry farmer". This cartoon illustrates two things: firstly, the NIMBYs (Not In My Back Yard) group of people; those who oppose the building of new basestations and those who challenge the safety of mobile phone use etc. Secondly the angry farmer illustrates literal land owners who Vodafone must pay rent to for mast space. The latter point becomes particularly pertinent when it comes to changing existing equipment, as the landowners and their lawyers tend to demand more money if the number of antennas increases or if other obvious changes occur. These considerations can conflict with technical decisions since many proposed advanced radio techniques demand additional antennas or more complicated equipment.

At the bottom left a cluster of competitor operating companies are shown for the UK market. Each Operator is constantly struggling to stay ahead of the others.

At the bottom centre, the shareholders are represented. Their influence is felt as a constant pressure to drive down costs. This is in tension with the goals of customer satisfaction and the operation and maintenance of a high quality network.

The cloud at the bottom right represents the services that users access from their phones. This is important for a couple of reasons. Firstly, different services have different quality of service metrics, and put different demands on the network in the form of minimum latency and data rate targets. Customer satisfaction is directly related to these factors along with service availability and thus knowing the expected traffic on the network is very important. A mystic with a crystal ball is shown to indicate that the operator is constantly trying to predict the future in terms of where the traffic is going and what is going to be the next big thing. The type of traffic that actually occurs determines the technical approach that must be taken, and particularly, the scheduling strategy.

Secondly, with the growth of dominant and pervasive online services, the operator is wary that it will become just a “dumb” bitpipe; simply another utility. The current marketplace can be likened to a real market in which the customer pays to enter, and the traders (Google, Youtube etc) get to setup their shops for free. This means that the Youtubes and Facebooks of the world, get free access to customers and the operator currently takes no share of the profits. Vodafone is aware of the economical advantage to being part of this market, either by offering differentiated quality of service to the highest bidder, or by offering its own services and applications. An example of this model can be seen in Amazon’s Kindle device, whereby the customer pays no subscription fee for book downloads, but instead the cost is included in the price of the book and the connection fee is charged directly to Amazon. These kinds of ideas have a direct technical trickle-down, since the system must be setup for service differentiation, which traditionally has not been the case.

Finally, and perhaps most importantly, the customer is represented in the top right of the diagram. Meeting the customer requirements is a central idea in Systems Engineering, and some believe it is *the* central idea. For example in [16], meeting the customer requirements is seen as synonymous with the central tenet of the book: building *quality* systems. Customer happiness is a major stakeholder in almost every technical choice. To take an example, the desire for good user throughput affects bandwidth decisions and thus is dependent on spectrum licensing and auction outcomes, it affects scheduling design and QoS (Quality of Service) stratification decisions, it influences the adoption of advanced antenna technology, and it motivates further research into interference mitigation, among other things.

1.2 Satisfying User Requirements

The last section examined the myriad and conflicting constraints placed upon the mobile operator when designing systems for the future. This section takes a closer look at the relationship between specific customer requirements and technical aspects of the network.

To illustrate this broadly, Figure 1-3 shows the first stage of a QFD (Quality Function Deployment), house of quality [17, 16], built using the free templates available from [18].

The methodology behind a QFD will now be described, followed with a discussion of the instantiation shown in Figure 1-3. The purpose of a QFD is to map soft, and often ambigu-

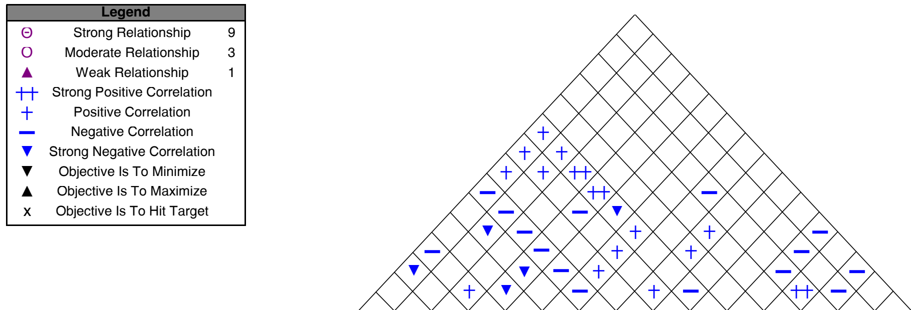

Row #	Max Relationship Value in Row	Relative Weight	Weight / Importance	Column #	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
				Direction of Improvement: Minimize (▼), Maximize (▲), or Target (x)	▼	▲	▲	x	x	▲	▼	x	x	▲	x	x	x	x	▲
				Quality Characteristics (a.k.a "Functional Requirements" or "Hows") / Demanded Quality (a.k.a "Customer Requirements", or "Wants")	Low order MCS scheme	High BW	High order MCS scheme	Scheduler that favours high SINR UEs	Scheduler that favours nobody	Frequency Reuse / ICIC	Simple components	Robust handover mechanisms	Scheduler that distributes interference	Scheduler that gives good instantaneous rates	Deals with top hardware vendors	Sufficent testing	Deals with service vendors	Deals with addon companies, cinema's	Good fixed line holdings
1	9	12.1	8.0	Low error rate	Θ	▲				▲		▲	▲						
2	9	12.1	8.0	High data rate		Θ	Θ	Θ		O				O					
3	9	7.6	5.0	Position independent experience	O	O	▲		Θ	Θ			O						
4	9	10.6	7.0	Cheap cost							Θ								
5	9	12.1	8.0	Operates on the move								Θ	Θ						
6	9	6.1	4.0	Low latency				O						Θ					
7	9	10.6	7.0	Fashionable phones											Θ				
8	9	10.6	7.0	Good services and applications												Θ	Θ		
9	9	10.6	7.0	Fashionable addons														Θ	
10	9	7.6	5.0	All in one packages / Tariffs															Θ
				Target or Limit Value	BW limited	Spectrum limited	SINR Limited	MaxCI	PF	Spectrum limited	Existing architecture	No handover failures	Standards limited	MCS limited	Monopoly limit	Cost limit	Monopoly limit	Monopoly limit	Monopoly limit
				Difficulty (0=Easy ,10=Extremely Difficult)	1	8	6	0	3	8	7	6	8	4	5	5	5	5	7
				Max Relationship Value in Column	9	9	9	9	9	3	9	9	9	9	9	9	9	9	9
				Weight / Importance	131.8	143.9	116.7	127.3	68.2	34.8	95.5	121.2	143.9	90.9	95.5	95.5	95.5	95.5	68.2
				Relative Weight	8.6	9.4	7.7	8.3	4.5	2.3	6.3	8.0	9.4	6.0	6.3	6.3	6.3	6.3	4.5

Figure 1-3: Top level QFD for the "network of the future"

ous, customer requirements, into tangible engineering targets such as program elements and physical features of products.

The first step in creating such a map, is to determine what the customer requirements actually are. The requirements are solicited through interaction with the customer, and various ethnographic tools are available for this purpose, for example: conversation, survey, participant observation, and so on. In the case of Figure 1-3 the customer requirements stated reflect the convictions of the author of this dissertation, in being an avid user of mobile phone products, as well as informal conversations with industry experts.

This informal approach is appropriate here, since the QFD is being used to reflect generally on the mobile communications industry and this dissertation's place in it, rather than systematically employing the QFD to develop a particular and specific product. Note that in the latter case, the outcomes of a top-level QFD would usually be exploded to form the inputs of further QFDs in order to tease out the necessary detail required for implementation.

Regardless, the process followed is not fundamentally different in either case. To map from customer to functional requirements, the customer requirements are first listed along the left of the QFD matrix as in Figure 1-3 under the heading "Demanded Quality".

The column to the immediate left of that column provides values indicating the importance of each requirement relative to the others listed. The chosen values are in this case subjective, but could in principle be empirically supported by survey data. The values are expressed as normalised percentages in the next leftmost column. They are used in a latter step, as will be explained.

Once the customer requirements are listed along the left of the QFD matrix, the next step is to list the quality (functional) requirements along the top. The quality characteristics are measurable characteristics of the end product [19]. That is, characteristics to which applied changes will directly impact the customer requirements. The characteristics are constructed by brainstorming for the customer requirements what tangible aspects of the network affect them.

Directly above each quality characteristic, a symbol is provided which indicates in which "direction" the characteristic would have to move to result in an improved product. For example, Column 7 lists "Simple components" and the symbol above is a down arrow, implying that greater simplicity will result in a better product (typically less failures, cheaper cost etc).

Below the body of the QFD, a target or limit value is provided for each quality characteristic. For example, for column 7 "Simple components" the limit value is "Existing architecture" implying that we do not desire to replace existing network components with simpler ones, so that in affect the aim is to minimise the addition of complexity to the existing network. Below this item, the difficulty of achieving each target or limit is estimated.

Once the quality characteristics are listed, the next step is to fill in the body of the QFD matrix. This part quantifies the relationship between the customer requirements and the functional requirements. For each pair in the Cartesian product of the two lists, a symbol is chosen to represent the strength of the relationship between them, as shown in the legend. Here a relationship is either non-existent (as indicated by no symbol), weak, moderate, or strong.

To give a concrete example of what has been explained so far, the customer requirement in Row 4 of "Cheap cost" is indicated to have a strong relationship with the functional requirement in Column 7 of "Simple components". This reflects the intuitive notion that simpler components are usually cheaper to make (this relationship evidently may not exist for all domains). The QFD says for this relationship: "the simpler the components used, the cheaper the end product will be". This mapping provides a useful overview of inter-relationships.

Following construction of the QFD matrix, the next step is to construct the "roof" of the QFD (the top triangle). The roof indicates the correlations between pairs of functional requirements.

The correlations are indicated by symbols for each pair of functional requirements. Each pair can be strongly or normally correlated in either a negative or positive direction. The absence of a symbol implies that there is no significant correlation between the given requirements.

Understanding these correlations is important for designing products as areas of mutual opposition imply the need for design compromises, whereas positive correlations point at areas for design simplification.

The final items in the QFD are the importance measures for each quality characteristic, these are computed as the sum of row products for each column. For each row in the given column, the row product is the product of the relationship symbol value and the relative importance of that row. For example, column one has a strong relationship in row one which has a value of 9 and the relative importance of row one is 12.1, therefore the row product is $9 * 12.1 = 108.9$, and the row product for row 3 is $7.6 * 3 = 22.8$. Therefore the sum of the row products is $108.9 + 22.8 = 131.8$ (when the non-rounded numbers are used in the spreadsheet). This value 131.8 is computed weight of the column. The final weighting therefore indicates how important each quality characteristic is because the value is proportional to how much that characteristic is related to important customer requirements.

Now that the QFD has been explained, we can broadly discuss it without any further clarification.

Examining the customer requirements: low user error rate, high user data rate, and mobility, are major (and obvious) customer requirements. Following this, cheap cost, fashionable phones, good extras, and single price plans are all high on the demand list. Other technical aspects, such as position independent experience and low latency, score low on the customer's requirements.

From the marketing perspective, in the past few years, fashionable phones, good applications, services, and tariffs have received considerable promotion and have been used as differentiating characteristics among operators. But this all changed considerably, shortly after the first iPhone was released. In London especially, O2 customers began complaining about data access problems and slow downloads. O2 attributed the problem to an 18-fold increase in data traffic, and had to spend 30 million on 200 new basestations to correct the problem [20]. Shortly afterward almost all operators suspended "Unlimited" data tariffs.

With this insight, and given the growing nature of mobile traffic, this problem is likely to worsen significantly in the future, even with the advent of LTE, to the extent that network quality and data rates will become important sales differentiators.

The first ten quality characteristics are technical, and the last more concerned with marketing and brokering good deals with suppliers. The work done in R&D (Research and Development) is technical in nature, and thus it is only possible to satisfactorily influence the first ten quality characteristics.

Two of these are concerned with MCS (Modulation and Coding Scheme) selection: lower order MCS schemes are needed for robust transmissions in low SINR situations, and higher order MCS schemes are needed to obtain high data rates under suitable SINR conditions.

Both these situations, and all intermediates, occur in the network, and so it should be clear that selection of the MCS set is an important aspect of network design.

Four of the quality characteristics are directly related to scheduling, and so it goes without saying that scheduler design is extremely important. This is because the scheduler has the final say over how resources are distributed to users, and can be setup to bias allocation toward particular sets of users if needed.

Low user error rate, high user data rate, and position independent experience, have a lot of correlations. All three of them are correlated are improved by increasing the cell bandwidth, since more resources are available for the same number of users. Low error rate and position independent experience also correlate with frequency reuse, in that increasing frequency reuse typically improves SINR conditions, and thus reduces the likelihood of error across the cell. Bandwidth and frequency reuse are however opposed, and are related by a strongly negative correlation. This is because as frequency reuse is increased, the network bandwidth must be divided bandwidth must be reduced.

High user data rate is strongly correlated with a scheduler that favours high SINR users, yet this is strongly negatively correlated with a scheduler that favours nobody, which in turn is strongly positively correlated with a position independent experience. Thus it is difficult to ensure high user data rates, and a consistent experience across the cell. This is because cell-edge users tend to consume a lot more resources per bit than do cell-centre users, and thus resources that could otherwise be used to give more users higher throughputs must be sacrificed to serve cell-edge users. There is a similar conflict with the scheduler that favours high instantaneous rates.

The cheap cost requirement is typically a function of network complexity and the complexity of the hardware employed. Simplicity is in conflict with almost every other technical quality characteristic. The operator must therefore act very carefully when it comes to adopting more advanced radio resource management techniques and multi-antenna solutions. The operator should be sure that gains can not be obtained through cheaper techniques applied to existing hardware.

This dissertation focuses on the third quality: position independent experience, and focuses first on the gains that can be leveraged from frequency selective scheduling to this end (Chapter 3), and then on the use of ICIC (Inter-Cell Interference Coordination) as a solution to improve the position independent experience (Chapter 4 and Chapter 5).

The next section describes in more detail each of these pieces of work, and relates them to the ongoing dynamic operations of the Vodafone business.

1.3 Business Context and Summary of Contributions

The three main contributions of this dissertation appear in Chapters 3, 4, and 5. The purpose of this section is to understand how these contributions relate to each other, and what motivated them, given the industrial context they were born out of. This is achieved by explaining in chronological order, how each new piece of work grew out of the answers provided by the last, modulated by the industrial drivers at the time. To this end, this section also summarises the methodology and outcomes of each contribution.

Section 1.3.1 discusses how the examination of frequency selective scheduling in WiMAX (Worldwide Interoperability for Microwave Access) seen in Chapter 3 was motivated by a technology "war" and a greater need to understand the competing systems. Section 1.3.2 explains how subsequent convergence on LTE as *the* next generation technology within the company, drove the need to find solutions to expected problems within LTE. And why interference mitigation was seen as important, causing it to become the object of analysis in Chapter 4. Finally, Section 1.3.3 explains how the results from Chapter 4, interpreted against the contemporary industrial context, lead to the focus on and development of dynamic interference mitigation schemes in Chapter 5.

1.3.1 Report 1: "*The Impact of MS Velocity on the Performance of Frequency Selective Scheduling in IEEE 802.16e Mobile WiMAX*"

When the RE (Research Engineer (RE) is the term given to a student undertaking the EPSRC (Engineering and Physical Sciences Research Council) programme for Doctor of Engineering) joined Vodafone Group R&D, the company was looking at radio network technologies for its next generation network, that is, the successor to 3G (Third Generation) and HSPA. OFDMA was considered the best candidate among the underlying transmission technologies on offer. Two OFDMA technologies were being considered: LTE and

Mobile WiMAX. The company was interested in understanding the technical limitations and performance of both systems. At the time the RE was appointed, the Mobile WiMAX standard was more mature than the LTE standard, and had better simulation capabilities within Vodafone R&D. It was believed that WiMAX and LTE, both being OFDMA systems, would to a first approximation, produce transferable simulation outcomes. Simulation of WiMAX was therefore considered a natural candidate for exploration.

The RE was tasked to work on a series of problems known collectively as SON (Self-Organising Network functionality). The concept of SON is that the manual legacy associated with previous technology deployments should be eliminated from forthcoming technology deployments. Subsequent maintenance and performance-related network modifications should be executed automatically in response to measured key performance indicators.

To this end, the RE became involved in an existing R&D collaboration called MONOTAS (Mobile Network Optimisation Through Advanced Simulation) [21], whose goal was to deliver "autonomous, self-optimising networks that are capable of reacting to rapid, unpredictable changes in the network traffic". This involvement motivated the research question addressed by the report in Chapter 3.

The research question came about as follows. With the advent of OFDMA came, for the first time (in practice), opportunistic frequency selective scheduling. It was not clear within the industry, what performance gains the paradigm would enable, nor how the scheme would perform under different mobility conditions. Due to a paucity of relevant research material, the company approved a proposal from the RE to investigate the relationship between the performance of frequency selective scheduling and MS (Mobile Station) velocity under 802.16e Mobile WiMAX.

Frequency selective scheduling was shown to improve performance by 100% under ideal conditions at 3km/h relative to frequency spreading. This performance was found to degrade linearly with increasing velocity until between 14km/h and 17km/h, the performance became worse than that of frequency spreading. The practical relevance of this research, given that the majority of users are projected to be stationary or have low velocities, is that frequency selective scheduling should be used by default in the majority of cells as opposed to frequency diversity. A 10% degradation in performance was observed when frequency selective scheduling was used at high velocities relative to frequency spreading, implying that measures should be taken to identify high velocity users and then use frequency spreading for them accordingly.

The work also discussed the probable impact of other WiMAX parameters on the performance of frequency selective scheduling, and speculated on the relative performance of LTE under similar conditions, before the LTE standard was finished.

The work was accepted for presentation at the Consumer Communications and Networking Conference in Las Vegas, January 2010 and appears in the IEEE (Institute of Electrical and Electronics Engineers) conference proceedings [22].

1.3.2 Report 2: "*The Impact of Scheduler Choices on the Performance of Static Intercell Interference Mechanisms*"

Shortly after completing the above work, the company made a technology decision and decided to focus on 3GPP LTE. Internal simulation capabilities for LTE had also improved by this point in concord with the first releases of the LTE standard. Subsequently, the RE became involved in a project called SOCRATES (Self-Optimisation and self-ConfiguRATion in wirelEss networkS)) [23] whose goal was to "effectuate substantial operational expenditure (OPEX) reductions by diminishing human involvement in network operational tasks, while optimising network efficiency and service quality". The RE was assigned to work on SOCRATES use case 3.1.1 [24] "Interference coordination" whose main objectives were to "Minimise the impact of inter-cell interference by managing the resources used in neighbouring cells" and "Maintain a fair balance between cell edge user performance and performance of users closer to eNodeB".

Some of the groundwork for the results presented in Chapter 4 was completed under the SOCRATES project and appears in the internal SOCRATES deliverable 3.1, whereupon the SOCRATES project entered a new phase of integration. Following this, the RE left the SOCRATES project and completed the majority of the work seen in Chapter 4 and all of the work in Chapter 5 was on behalf of Vodafone R&D for ongoing internal programmes outside of SOCRATES.

As this was going on, the company was trying to decide what features and claims, made by equipment vendors of the new LTE technology, were likely to stand up in practice. Reports were being received by the company from equipment vendors, pushing the concept of soft frequency reuse and making positive claims about it. It was therefore considered prudent to subject these claims to a critical assessment.

Chapter 4 provides a comprehensive analysis of soft reuse in LTE. The factors which contrive to determine the success or failure of a static reuse scheme are identified as: SINR distribution shift under application of technique, SINR to MCS lookup curve shape, and proportional use of different MCS schemes due to UE (User Equipment) distribution and scheduling strategy. The interaction of these factors is illustrated with a quantitative example.

The work goes on to challenge the common assumption that it is "best" to give resources with a high reuse factor to those at the cell-edge, by showing for a fixed rate service class, that it is best to be greedy and give these resources to those at the cell-centre. The example serves mainly to illustrate the potency of the interaction between interference environment, scheduler, and performance metric.

The final section of the work addresses the interaction of soft reuse and scheduler more generally, by enforcing different degrees of fairness upon the scheduler, and assessing the performance of soft reuse relative to reuse one under these conditions. The work was performed using monte-carlo simulations, in the downlink direction, on a London scenario with realistic path loss and network data. All work was statistically quantified using appropriate tests.

The final conclusion was that soft reuse can be tuned, via varying the ratio of power on the different resource bands, to provide optimal performance for a given user distribution, SINR distribution, and MCS lookup curve. More generally, it is observed that the optimal operating parameters are *specific to* particular cell configurations and that no generalised gains can be obtained. Given the heterogeneity of extant and future deployments, the discovery that no soft reuse scheme tested provides a generalised gain independent of these factors, leads to the recommendation that static interference avoidance schemes be abandoned in favour of dynamic ones, that can adapt to changing conditions.

1.3.3 Report 3: "*Intrasite Scheduling for Interference Avoidance in LTE*"

Following on from the solid foundation provided by the static scheme analysis, and concluding that nothing suitable for a heterogeneous network can be found there, it was decided that the RE should focus his efforts on dynamic schemes which adapt to fast fading and other time sensitive changes. This was accompanied by a company decision to deploy LTE networks under reuse one (where cell shares the same bandwidth), and whilst this does not prevent static patterns being created, through careful control of transmit powers, it is a clear indicator that the best way to proceed is via dynamic schemes that operate within a reuse one environment.

There is a broad spectrum of work on dynamic ICIC schemes, some of which is interesting and offers promising performance gains. A purely practical weakness met by the majority of this work however, is caused by the economic environment into which cellular network technology is deployed. In many markets, a number of different vendors supply equipment to the network operators, which means that any non-standardised distributed algorithm would have to be supported across multiple hardware environments and by multiple vendors. Whilst the current LTE standards specify some communication principles with which ICIC can be carried out, *they do not presently specify specific algorithms for ICIC*. Arguably, this is unlikely ever to happen, since vendors have historically differentiated themselves on the proprietary aspects of their equipment. This means that even many interesting and technically feasible schemes proposed in the literature, have a low likelihood of use in practice, more especially so since vendors will tend to leverage their own IPR (Intellectual PRoperty) in favour of purchasing new ideas. In this sense a lot of work completed to date serves best to inform decisions at a more abstract level, rather than to dictate particular solutions.

Given this insight about the market, it was decided to focus on solutions with a higher propensity to lead to practically useful results. To this end it was decided to focus on intrasite coordination. The simplest form of inter-cell coordination is intrasite coordination. It is simple because the baseband equipment for the sectors at a given site is usually co-located. Meaning that the entry barrier to inter-cell communication is low, which increases the likelihood of adoption. In addition, since a given site is usually under the control of only one vendor, the need for multi-vendor support is reduced if not mitigated. The operator need only convince one vendor to implement the promising solution, or at least enable it, instead of having to convince *all* the vendors to cooperate in supporting the idea. In some cases, where the operator in question has a substantial majority position, they may even be able to explicitly request the inclusion of a given feature.

In Chapter 5 an algorithm is developed for intrasite interference coordination that produces robust albeit modest performance gains, when compared with sector-independent frequency selective scheduling. The gains are evident whenever the cells at the site are not fully loaded, and whenever the scheduling strategy is not overly biased toward low SINR UEs.

With reference to market considerations, the results were influential in a decision to mandate the ability for intrasite communication to a vendor, with the requirement appearing in the specifications.

The work also lead to the identification of a problem with CQI (Channel Quality Index) estimation: namely that the conditions at the points of measurement and use of CQI can be substantially different. A solution to this problem is being developed internally as a patent submission, leveraging the information that the eNB *knows* about intrasite resource usage at the time CQI measurements were made.

1.4 Chapter Summary and Outline of Dissertation

Section 1.1 argued that the mobile phone industry is a complex multi-organisation *system* that is pulled in multiple directions by the sometimes competing goals of a multitude of incumbent stakeholders. Vodafone's most important stakeholder was identified as the customer. Vodafone's customers were shown to place diverse, and multi-faceted, demands upon the network, and to introduce equipment with disparate capabilities and operating requirements. The challenge faced by Vodafone was identified as being to satisfy as many customers as possible, whilst minimising the impact on other emergent properties of the system. Toward this goal, Section 1.2 identified correlations between the technical aspects of the network and the satisfaction of customer requirements. This motivated the desire to use scheduling and interference mitigation solutions as a vehicle for network improvement.

In Section 1.3, the industrial context and key drivers were identified for each of the three main contributions of this dissertation (described in Chapters 3, 4, and 5). How each contribution lead on to the next was explained, and the outcome of each contribution was sketched.

In summary, the primary outcomes of this dissertation are: (i) a novel dynamic intrasite interference mitigation technique (ii) the ongoing internal development of a patent for improving CQI estimation to this end. The latter complements the former, as well as providing more general improvements in opportunistic scheduling performance. There are two secondary outcomes: (i) link and system level analysis of opportunistic scheduling gains in 802.16e Mobile WiMAX, (ii) system level analysis of static inter-cell interference strategies in 3GPP LTE and derivation of determinant performance factors.

The rest of this dissertation is organised as follows: Chapter 2 reviews the relevant literature for each of the research topics covered, and provides some general technical background for non-experts. Chapter 3, Chapter 4, and Chapter 5 contain the three main research contributions. Chapter 6 provides some preliminary analysis concerning extension of this work and Chapter 7 concludes the dissertation and discusses future work.

Disclosure

The reports in Chapter 3, Chapter 4, and Chapter 5 appear verbatim as they were published internally within Vodafone R&D. There is some level of duplication within the dissertation because of this. A reduced version of the work in Chapter 3 appears publicly in a conference paper [22]. A reduced version of the work in Chapter 4 appears publicly in a conference paper [25]. The full list of publication outcomes is shown below:

- Ashley Mills, David Lister, Marina De Vos and Yusheng Ji. The Impact of MS velocity on The Performance of Frequency Selective Scheduling in IEEE 802.16e Mobile WiMAX. In *Proceedings of the 7th Annual IEEE Consumer Communications & Networking Conference (IEEE CCNC 2010)*, Las Vegas, Nevada USA, January 2010.

- Ashley Mills and David Lister. Intrasite Scheduling for Interference Avoidance in LTE. In *To appear in proceedings of: 2011 IEEE 73rd Vehicular Technology Conference: VTC2011-Spring, 15-18 May 2011, Budapest, Hungary*, 2011.

- Ashley Mills and David Lister and Marina De Vos, "Understanding Static Inter-Cell Interference Coordination Mechanisms in LTE," *In press (available upon request).*

Chapter 2

Literature Review

Whilst the reports in Chapters 3, 4, and 5 contain background material and literature references to motivate them, the depth to which this is executed falls short of demands associated with academic publication. The reason for this is that the industrial audience for whom the reports were written, required a different style of presentation. However, for the purpose of this dissertation, an in-depth review of the literature is necessary, in order to rigorously establish the grounds for the research work that follows.

This literature review is split into three sections. There are two principal reasons for this. Firstly, in order to make this dissertation more accessible, a technical introduction to cellular networks has been provided. This background is common to the rest of the work and so to avoid duplication, appears only once, here in Section 2.1.

Secondly, the research work of this dissertation originated from two different time periods within the company, under slightly different motivational contexts, and concerns different OFDMA network technologies. Whilst broadly speaking the same goals are sought, as was outlined in the introduction, the technology differences and approaches taken are different enough to warrant separation into two parts. For this reason Section 2.2 provides the background and literature review for the work on frequency selective scheduling in WiMAX, whereas Section 2.3 covers the work on static and dynamic interference mitigation strategies in LTE. The structure of each section is as follows:

Section 2.1 provides an overview of the most important technical elements found in cellular networks and introduces the terms and concepts that occur later in the reports. This information is provided to make this dissertation accessible to a wider audience and to disambiguate nomenclature. This may be skipped by technology experts without loss.

Section 2.2 provides background information for the first report (Chapter 3). It contains an introduction to the relevant aspects of 802.16e WiMAX. In addition to providing a overview of the physical layer, the different scheduling "zones" that WiMAX enables are described. The scheduling zones allow frequency selective scheduling and frequency diverse scheduling to be used, which are relevant because they are the objects of comparison in Chapter 3.

Section 2.3 provides a comprehensive literature review of static and dynamic interference avoidance schemes, and describes how this literature motivated and informed the work presented in Chapters 4 and 5.

2.1 General Technical Background

This section serves to introduce non-experts to the concepts and terms that are used in the cellular networks research literature and will allow non-experts to process the rest of this dissertation. Readers with prior knowledge of cellular network terminology and operation may wish to skip this chapter, although it may still prove useful to skim it with a view to disambiguating nomenclature.

2.1.1 Cells

Cellular networks provide support for bi-directional communication between mobile and spatially separated UEs. This is achieved by dividing the desired coverage region into a multitude of tessellating hexagonal areas called "cells". Each cell contains a transponder called a basestation (BS). The BSs are connected to a wired backbone network which enables them a permanent connection with each other. UEs can communicate with each other over this backbone by establishing wireless connections to their nearest BSs. This concept is illustrated in Figure 2-1.

The figure is very simplistic and contains several abstractions; the wired connection shown

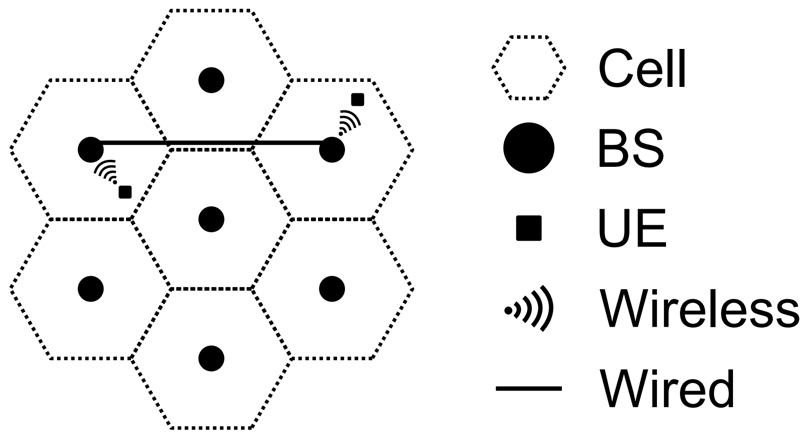

Figure 2-1: Communication in a cellular network. Two UEs are shown communicating with each other via their nearest BSs. Their connection is routed between BSs via a wired backbone connection.

represents an intermediate network, which may itself contain further wireless links such as direct microwave connections. The point is that the intermediate wired network is permanent, and in contrast to the UE $\rightarrow$ BS link, immobile.

Furthermore, the BS network contains in practice additional layers of abstraction and higher order management nodes. Real networks contain databases to keep track of UE locations, encryption principles to protect conversations, accounting protocols to keep track of bills, routing nodes to control traffic and congestion, gateways into different services such as The Internet, and other things. These latter functions are implementation specific and will be introduced later if relevant to the discussion. Suffice it to say for now, that the network is endowed with everything required to setup, support, and terminate connections between any two UEs that are registered to it.

UEs are mobile and can move across cell boundaries. When this happens, the BS with which the UE communicates, called the serving BS, is changed to the BS in the new cell. The process of changing the serving BS is called handover.

2.1.2 Sectorization and frequency reuse

The cellular concept was conceived to allow mobility over large geographical areas. In Figure 2-1 each cell contains a single BS, implying a single omni-directional antenna at the centre of each cell. Whilst this is technically more efficient, the acquisition of sites for BSs is very costly. And so for economical reasons, to increase the capacity per site, the circumference around each antenna site is usually split into spatially disjoint areas in a process called sectorization. This is illustrated in Figure 2-2.

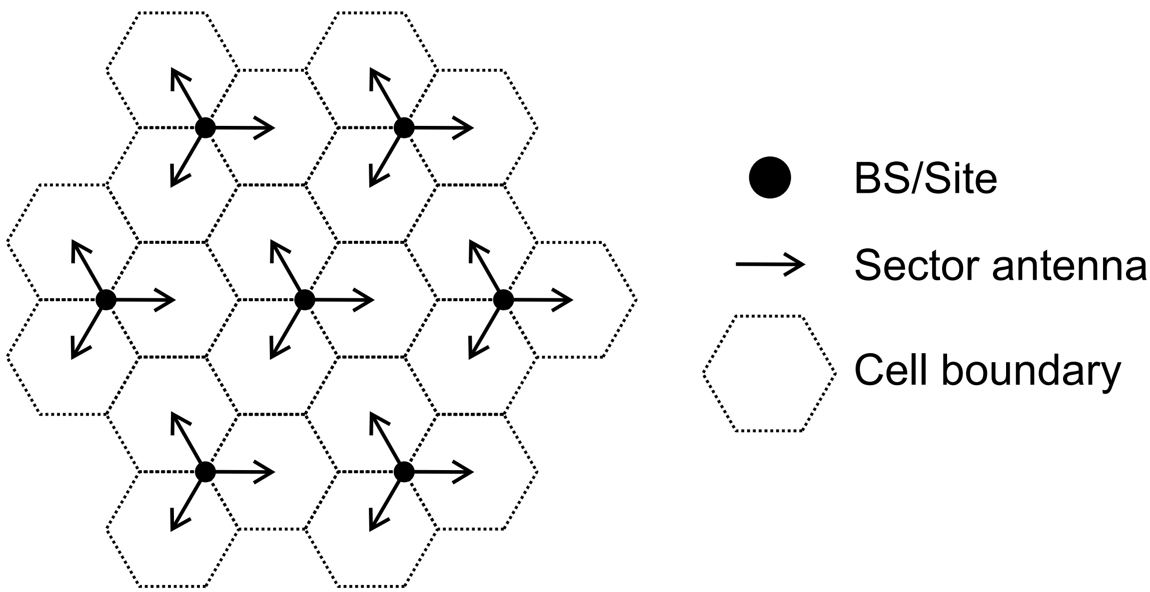

Figure 2-2: The circumference of a BS at one site is usually sectorized into multiple antennas, each of which is directed to serve a separate cell.

In the figure, each site has three antennas. Each antenna is setup to radiate its signal over a spatially disjoint geographical area. The terms cell and sector are equivalent and are used interchangeably in this document.

When a BS transmits or receives, the information is carried over spectrum licensed to the network operator. This is finite in quantity, and so it is not possible for each cell to transmit using a different frequency and at the same time have usable bandwidth.

It is inevitable therefore that spectrum will be reused. This concept is called frequency reuse, and is characterised by a so-called frequency reuse factor, sometimes shortened to reuse factor.

The reuse factor specifies how many disjoint frequency sets are used by the network. For a reuse factor of f, the available bandwidth BW is divided into f disjoint regions $\{b_1, b_2, ..., b_f\}$, and each cell is assigned a single region b_i. Whenever the number of cells in the network is greater than f (the normal situation) cells will be forced to share b_i.

Since the cellular network is designed to have no geographical gaps in it, this means that cells using the same b_i will interfere with each other. There are two approaches to overcoming this problem:

1. Use a frequency reuse factor of one, such that all cells use the same frequency, and encode the data in such a way that interference can be tolerated.
2. Use a frequency reuse factor higher than one, and try to ensure that adjacent cells use different frequencies.

Both these approaches are used in deployed technology. The latter is illustrated in Figure 2-3, where a reuse factor of three is illustrated.

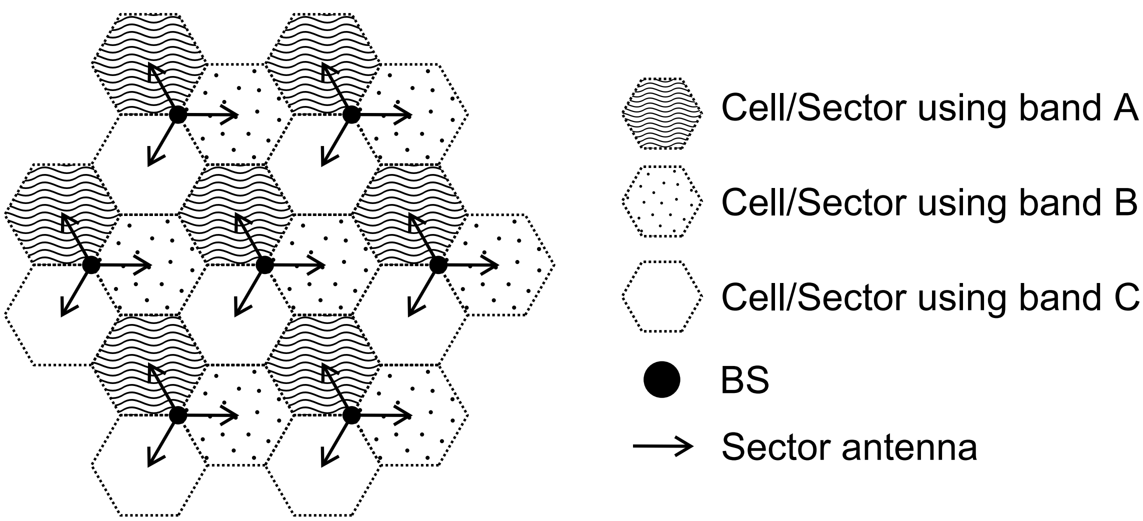

Figure 2-3: Frequency reuse with a reuse factor of three.

The reuse factor uses different frequency bands in adjacent cells such that interference is reduced. Note however that since the available bandwidth is split into three, each sector has less bandwidth available to it than if a reuse factor of one had been used. And so a tradeoff exists in reusing frequency between improving interference at the cost of reduced bandwidth. This tradeoff is the object of study in Chapter 4 where a specific combination of reuse one and three called soft-reuse is examined.

2.1.3 Antennas

In practice, antennas are not deployed on perfect hexagonal grids as was indicated in the last section; firstly, because sites can be deployed only where land and site licenses can legally be obtained, and secondly, because the Earth's surface is not a perfect geometric plane but is rather beset with various undulations and myriad changes in elevation. Any practical deployment will thus be an irregular approximation to its geometric analogue.

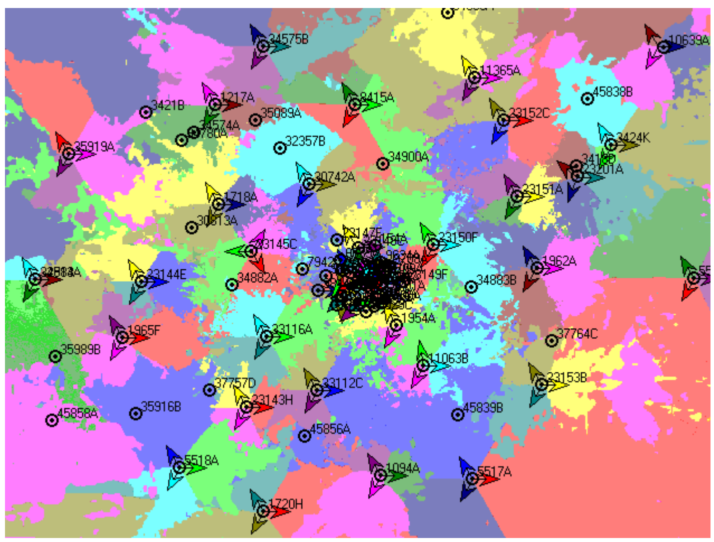

Figure 2-4: Antenna deployment in and around the city of Leeds.

Figure 2-4 shows an antenna deployment in and around the city of Leeds. A large density of antennas can be seen at the city centre and then surrounding this, where the population is less dense, the antennas fall into a more regular pattern. The deployment is hexagon-like as can be established by comparison with Figure 2-2. However, the shaded areas associated with each directional antenna are coverage predictions, and from this it should be clear that reality departs significantly from the idealised hexagonal shape.

The signal which emanates from an antenna is directed, anisotropic, and has a distinct 3D profile that is not a hexagon. Figure 2-5 shows such a 3D profile.

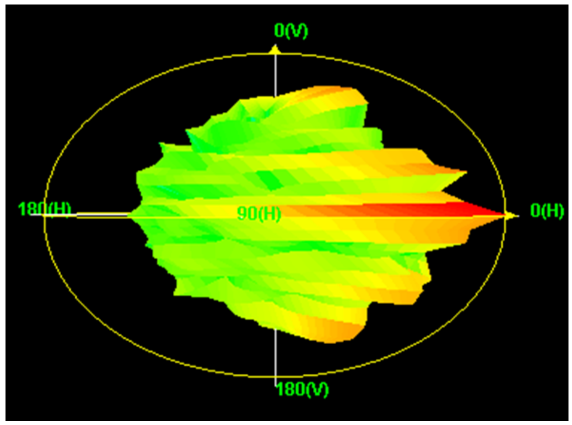

Figure 2-5: EMR pattern of a directional antenna in 3D.

Figure 2-6 shows the antenna pattern, as claimed by the specification sheet, for a Kathrein antenna. The antenna pattern is oval shaped when viewed from the horizontal, and is roughly sausage shaped when viewed from the vertical. The vertical pattern is able to be tilted up and down between 0° and 14° via a mechanism called electrical tilt.

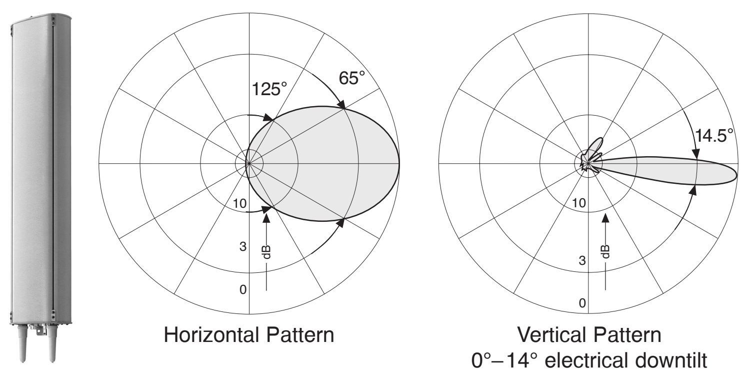

Figure 2-6: Antenna specification diagram. The antenna shown on the far left has a beam pattern which is approximately described by the two figures shown to its right.

The beam of an antenna covers a particular area for transmissions from BS to UE. The pattern thus indicates coverage in the downlink (DL) direction. However, channel conditions are usually assumed to be reciprocal and hence the area the beam covers can also be considered to cover the region of uplink (UL) communication for that antenna.

For simulation, is is common to ignore the vertical pattern and assume that the UE is within the horizontal plane. The horizontal pattern is usually specified analytically as in Equation 2.1

$$A(\theta) = -min\left[12 \cdot \left(\frac{\theta}{\theta_{3dB}}\right)^2, A_m\right] \tag{2.1}$$

Where A_m is the minimum gain in dB and θ_{3dB} determines the $3dB$ beamwidth, which is defined as the angle at which the attenuation from the boresight[1] is $3dB$. This is plotted in Figure 2-7 for $\theta_{3dB} = 70°$, $A_m = 20dB$.

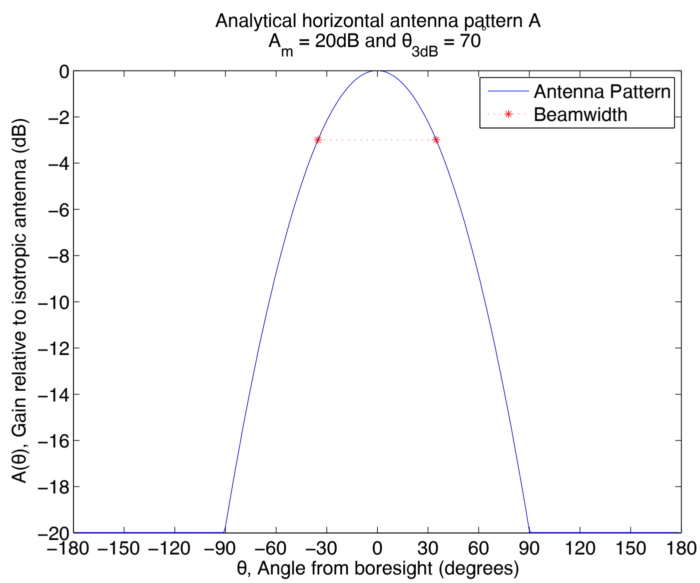

Figure 2-7: Typical analytical antenna pattern

To explain what this means, consider a theoretical isotropic antenna that radiates the same power in all directions. Directed antennas modify this pattern so that the power radiated is anisotropic. The pattern is thus specified in terms of signal gain relative to an isotropic antenna.

The pattern itself is usually normalised so that the main lobe (the boresight) of the antenna

[1]The boresight is direction at which the antenna transmits with most gain

pattern has a gain of $0dB$. This usually means that the rest of the pattern has a negative gain, as seen in Figure 2-7. The reason for doing this is so that the gain attributed to the main lobe can be controlled separately by an offset called the boresight gain.

For a boresight gain of xdB one would expect to measure a signal xdB stronger than the transmit power of the antenna at the boresight.

BSs are deployed at different heights depending on site availability and the surrounding ground will vary in elevation. In order to cover a desired area, it is necessary therefore to adjust the properties which control the shape and direction of each antenna beam and by consequence, the area over which that beam provides coverage for UEs.

Cellular network antenna beam direction is controlled by two parameters as shown in Figure 2-8. The first parameter, Azimuth, is shown on the left for a tri-sectored basestation (3 beams), it controls how much the antenna is rotated clockwise by an angle $A1$ from magnetic north and thus determines the direction across the land over which the three beams travel. The second parameter, Downtilt, is shown on the right, it specifies how much the antenna is tilted clockwise by an angle $A2$ from the horizontal plane.

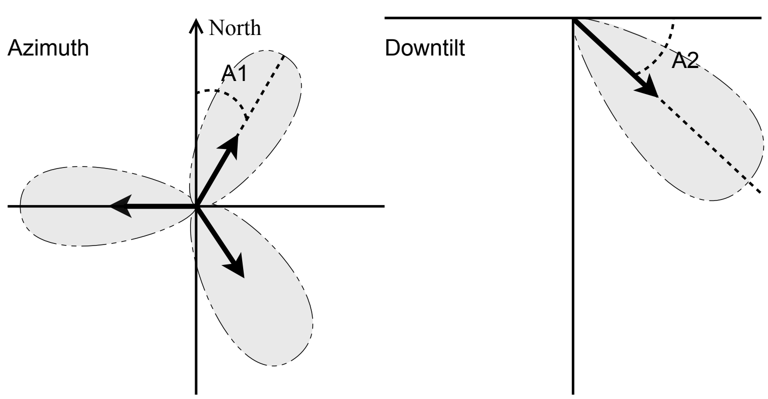

Figure 2-8: Antenna rotation and tilt adjustment. The concept of rotational adjustment (azimuth) is illustrated on the left whereas vertical tilt adjustment is illustrated on the right

Downtilt can be controlled either mechanically, where the antenna complex itself is physically tilted, or electrically, where the power is delivered to the antenna components in such a manner as to generate a beam shape at the horizontal which is comparable to that obtained by a mechanically downtilted antenna.

Thus downtilt is often specified as a combination of electrical and mechanical downtilt or uptilt, for example, an electrically tilted antenna may be mechanically uptilted by a few degrees to obtain the desired beam profile.

It is worth noting here that the hexagonal arrangement shown in Figure 2-2 is sometimes drawn as in the right side of Figure 2-9.

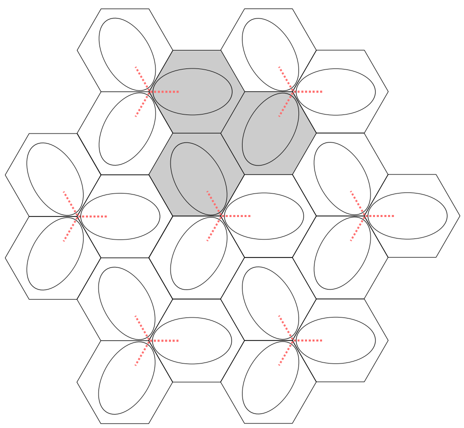

Physical layout

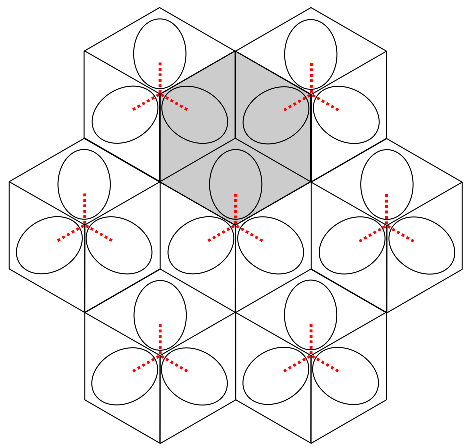

Logical layout

Figure 2-9: Tessellation of trisector base stations and a logical view of this.

This presentation is sometimes preferred when it is required to clearly delineate the regions of interference associated with a triad of converging sectors, for example in the description of algorithms. As can be seen, the region of interest, shown shaded, is logically represented on the right as three sectors whose beams converge, whereas physically, the interfering beams are interlocked in a triangular pattern as shown on the left.

2.1.4 Signal Propagation

Overview

In a vacuum, the received power P_r in dB relative to $1kW$ of a transverse wave transmitted from a point source with power P_t in the same units is as follows [26]:

$$P_r = P_t \left(\frac{\lambda}{4\pi d} \right)^2 \tag{2.2}$$

Where d is the distance between transmitter and receiver in km and λ is the wavelength of the transmitted signal. The ratio P_t/P_r is known as the transmission loss. This is often referred to as the pathloss and is normally expressed logarithmically as a positive attenuation:

$$PL = 32.44 + 20 \log f + 20 \log d \tag{2.3}$$

Where f is the transmission frequency in MHz and d is the distance in km.

Unfortunately, cellular communications do not occur between point sources in vacuums and are in reality subject to many kinds of physical attenuations and distortions peculiar to the chosen antenna systems and operating environment.

The basic pathloss equation above has been modified by researchers, based on empirical measurement, to try and reflect the characteristics of various environments and this is discussed in the next section.

In addition to pathloss, which can be seen as somehow capturing the most basic attenuation of the signal through space, there are many other factors that affect signal quality which can be categorised broadly as follows:

- Antenna pattern. As seen above, the antenna pattern controls how strong the signal is relative to the angle that the UE makes with the BS.

- Pathloss. As already described, this is concerned with the general inverse square decay principle associated with the transmission of electromagnetic radiation, and sometimes it is modified to account for general environmental trends.
- Slow fading. This is generally caused by large objects in the transmission path and causes large changes in signal attenuation over relatively slow time periods from seconds to minutes.
- Fast fading. This is generally caused by small objects in the transmission path and causes relatively smaller changes in signal attenuation but which occur over smaller time periods from nanoseconds to several hundred milliseconds.
- Noise. This concerns attenuation due to thermal noise, cable loss, and other intrinsic factors.
- Interference. This is attenuation due to other transmitting devices.

When modelling overall transmission loss, the resultant equation is typically the sum of all these factors. The next five sections will briefly review each type of transmission loss that contributes to the overall sum.

Empirically derived pathloss models

As previously mentioned. The generalised pathloss equation, Equation-2.3 as been modified for use in simulation and modelling, using observed data to account for various field conditions. There are many pathloss models [27], and they all take the same general form of Equation-2.3 above. For the sake of example, the Okumura-Hata Model and an extension of it, COST-231 will briefly be examined.

The Okumura-Hata propagation model is based on data derived in Japan by Okumura et al. in 1968. Hata later developed a mathematical formulation based on this data which became very influential in the field of cellular network simulation and planning [26]. The Okumura-Hata model has been extended at various times to account for frequencies above 1Ghz, and for different clutter scenarios [27]. In this work, the extended Okumura-Hata pathloss model as specified in [28] was used. This model was chosen based on recommendations in [29].

The model defines three environment types: URBAN, SUB-URBAN, and OPEN-SPACE. Here the name RURAL is substituted for OPEN-SPACE. The equations for the models are defined below:

$$\begin{aligned} \mathrm{URBAN}(d) = \quad & 46.3 + 33.9\log(f) + 10\log(f/2000) - 13.82\log(\max(30, H_b)) \\ & + [44.9 - 6.55\log(\max(30, H_b))]\log(d) - \alpha(H_m) - \beta(H_b) \end{aligned} \tag{2.4}$$

Where

$$\begin{aligned} \alpha(H_m) = \quad & (1.1\log(f) - 0.7)\min(10, H_m) \\ & -(1.56\log(f) - 0.8) + \max(0, 20\log(H_m/10)) \end{aligned} \tag{2.5}$$

and

$$\beta(H_b) = \min(0, 20\log(H_b/30)) \tag{2.6}$$

$$\mathrm{SUBURBAN}(d) = \mathrm{URBAN}(d) - 2\left\{\log(\min(\max(150, f), 2000)/28)\right\}^2 - 5.4 \tag{2.7}$$

$$\begin{aligned} \mathrm{RURAL}(d) = \quad & \mathrm{URBAN}(d) - 4.78\log(\min(\max(150, f), 2000))^2 \\ & +18.33\log(\min(\max(150, f), 2000)) - 40.94 \end{aligned} \tag{2.8}$$

Where H_b is the height of the BS antenna in m, H_m is the height of the mobile in m, f is the carrier frequency in Hz. And d is the distance from the antenna in m.

The output of the model used is shown in Figure 2-10 along with the free space propagation model defined in Equation 2.3 which is called Line of Sight (LOS) in the figure.

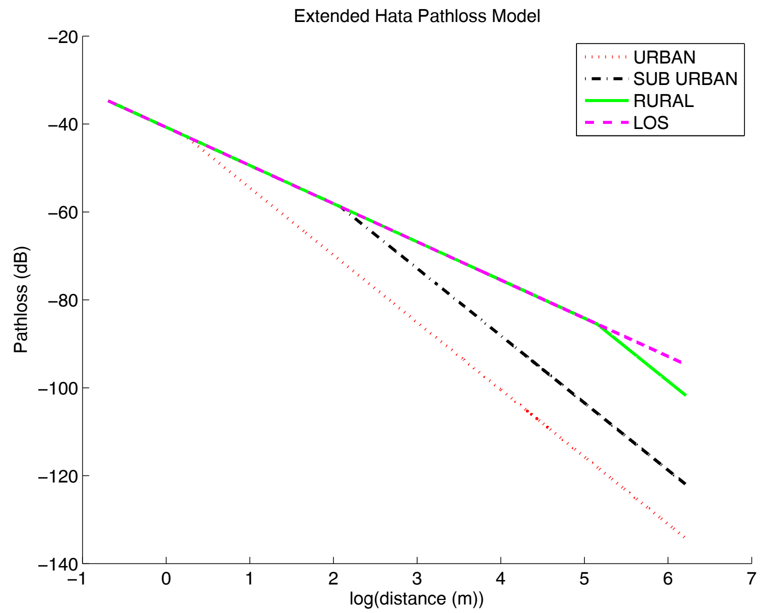

Figure 2-10: Various scenarios of the Extended Hata Pathloss model

Slow fading

Large scale fading is often called shadowing owing to a conceptual similarity with another form of electromagnetic radiation, namely light. In the same way that a person obscured from the sun by an object will experienced diminished light in the shadow cast by that object, although being not altogether be plunged into darkness, the same happens to a mobile user when a large object obscures its line of sight path to the base station with which it is communicating; the radio signal is diminished but not altogether eliminated. This is illustrated in Figure 2-11.

To capture the effects of buildings, the generalised pathloss equations given above are perturbed by a log normal random variable, which in decibels is simply drawn from a zero mean Gaussian. The term is simply added. The standard deviation of the Gaussian will usually be based on empirical data.

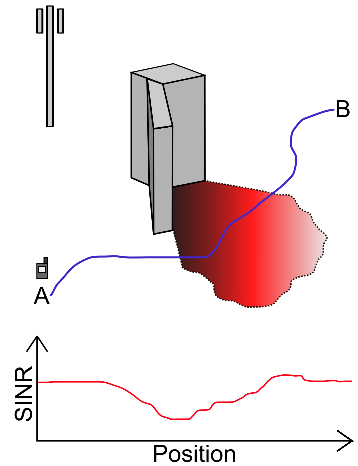

Figure 2-11: An illustration of a slow fade that appeals to intuition; the mobile user travels from point A to point B and in doing so is “shadowed” by the building. The plot of SINR vs distance travelled shown below the figure shows the correlated drop in SINR as the user passes through the radio shadow cast by the building.

Fast fading

The following factors contribute to fast fading:

- Absorption.

  Various objects, including atmospheric particles, absorb microwave radiation. Objects have different absorption spectrums, meaning that they absorb different wavelengths to different extents. This contributes to the generation of a received power profile that varies in amplitude across different frequencies.

- Diffraction.

  When a signal is diffracted, not only is it diluted, but the chance of later multipath effects is increased.

- Reflection and scattering.

  Microwaves signals reflect off certain surfaces, which can lead to the signal taking multiple paths to the receiver and consequent multipath reception.

- Multipath reception.

Mainly due to the action of diffraction and reflection, multipath reception is where a signal reflects off two or more surfaces and arrives at the receiver at different points in time. The received signals sum up to result in a received power profile that varies in amplitude across different frequencies. Multipath reception is summarised in Figure 2-12.

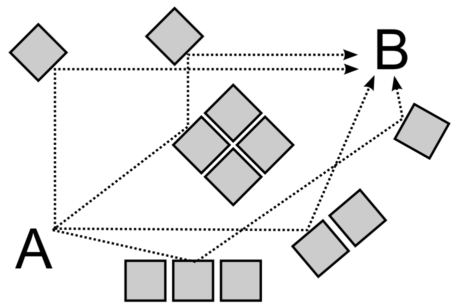

Figure 2-12: An intuitive illustration of multipath fading; the figure shows transmissions from a point A to a point B where there is no line of sight path available. Thus transmissions are achieved via reflections. Each reflection introduces a delayed version of the signal which undergoes frequency selective fading according to an individual profile. These reflections are combined at the receiver.

- Doppler shift.

  The relative motion of the UE to the BS can result in Doppler shifts of the signal, which means it will be received at either a higher or lower frequency than it was transmitted. Combined with multipath reception, this leads to a received power profile that varies in amplitude across different frequencies.

Fast fading results in a received power profile that varies in amplitude across different frequencies. This is called a frequency dependent power profile.

This profile is also dynamic, in that it itself changes through time. This is because, as the UE moves, so does its position relative to the BS and hence the set of objects in between the two which contribute to reflection, diffraction, and absorption will also change. Note that even a stationary UE may experience fast fading because the objects in between UE and BS may themselves be moving.

The basic approach to modelling fast fading is to assume the received signal is composed of several multipath components. This signal profile can be approximated by a Rayleigh distribution [26, 30].

A Rayleigh fading channel is modelled as the sum of several independent flat Rayleigh channels . Mathematically, such a sum can be written:

$$p(k) = \left| \sum_{p=1}^{paths} M_p A_p e^{j[\theta_p - 2\pi f_k T_p]} \right|^2 = \left| \sum_{p=1}^{paths} M_p A_p e^{j\theta_p} \cdot e^{-j2\pi f_k T_p} \right|^2$$

Where M_p is the amplitude vector of the discretely sampled input Rayleigh signal, A_p is the gain for path p, f_k is the kth frequency of interest, and T_p is the delay in seconds associated with path p.

Figure 2-13 illustrates the frequency and time dependent nature of the power at which a transmitted signal is received. Each column contains three plots; the rows correspond to UE velocities of respectively $1km/h$, $15km/h$, and $30km/h$.

The left column shows the frequency dependent nature of received power, each plot shows how the amplitude of the received signal varies with frequency. The left plots contain lines for $0ms$, $1ms$, and $2ms$ after transmission to illustrate that the signal also varies with time. The right plots focus on time variance and show how the received signal amplitude varies with time.

The right column of plots show this time variance for 3 subcarriers each to illustrate that the received signal also varies with frequency. As can be seen from either column, in general, the temporal variance increases as the UE velocity is increased. This is because the objects which determine this profile change more rapidly as UE velocity is increased.

There are many sophisticated models available, with characteristics similar to the model examined above, but with extensions for multiple antennas, and other effects. A basic Rayleigh channel is used in Chapter 3, and the Extended Spatial Channel Model (SCME) [31] is used in Chapter 4 and Chapter 5. These are both widely used in the literature, with the latter gaining much popularity of late.

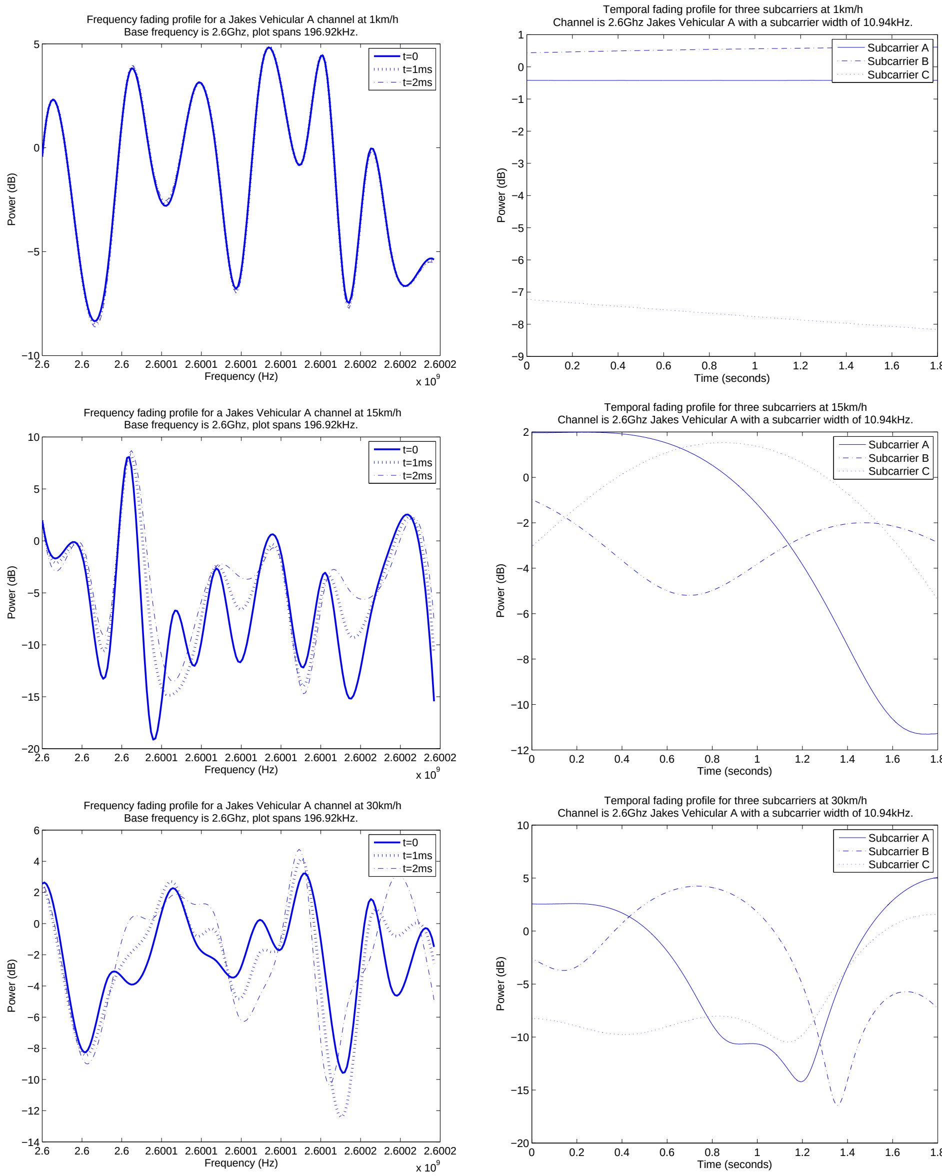

Figure 2-13: The frequency-time dependent nature of received power. The left column shows frequency vs received power for a Jakes Vehicular A channel at increasing UE velocity whereas the right column shows time vs received power at increasing UE velocity. In each left column plot the frequency response curve is shown at $0ms$, $1ms$, and $2ms$ after transmission to illustrate the variance of the frequency dependent response through time. In each right column plot, the time dependent power profile of three adjacent subcarriers is shown.

Noise

Thermal noise is present in all electrical systems. It is typically modelled as Additive White Gaussian Noise (AWGN) which is nothing other than perturbation by an independently drawn normally distributed random variable. Figure 2-14 shows the Sin function perturbed by AWGN.

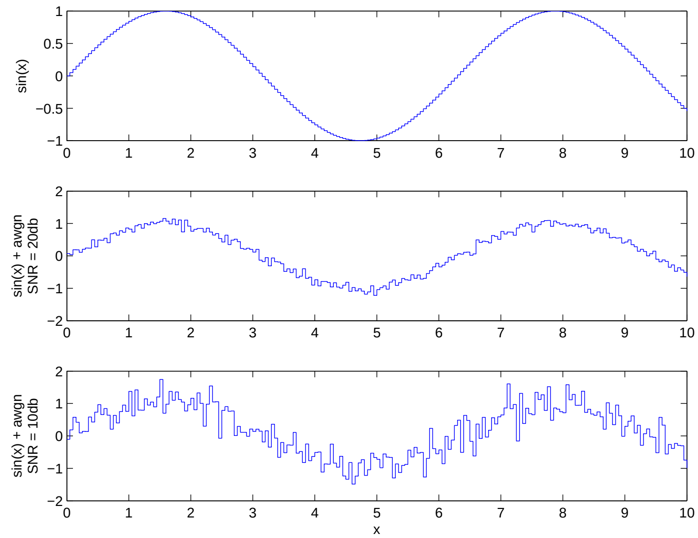

Figure 2-14: Sin wave corrupted by increasing amounts of white noise. From top to bottom: no noise; SNR of 20db; and SNR of 10db.

The first subplot shows a stair plot of a sine wave sampled at intervals of 0.05. In the second subplot the same sine wave is shown corrupted by an AWGN source. The AWGN source was generated so as to have an SNR relative to the sine wave of 20db. The third subplot shows the sine wave corrupted again by an AWGN source, this time with a relative SNR of 10db.

Interference

A large source of interference for cellular systems is interference from other cells. This is modelled in exactly the same way as a desired signal and must be computed for each transmission.

Other losses

Finally, losses can come from other sources that do not fit neatly into the other categories:

- Coupling loss due to mismatch of transmitter and receiver antenna orientations.
- Cabling and other internal equipment losses.
- The Sun and other natural causes also result in some general background interference.

2.1.5 Sharing Finite Resources (Multiplexing)

Introduction

Radio access technologies work given a fixed region of bandwidth within which to operate. To support more than one user, and bi-directional communication, the bandwidth must be split and shared.

First used in 1885[2], the general term for this is multiplexing. The Oxford dictionary definition is "to enable (a line) to carry several signals simultaneously" [**?**].

In contemporary telecommunications language, the term "duplexing" is usually reserved to refer to the splitting of a communications channel into uplink and downlink regions. Whereas the term "multiplexing" is used to refer to the sharing of a communication channel in a given direction into multiple logical channels.

Uplink and Downlink Channels

Downlink and uplink transmissions are separated in either time or frequency, as illustrated in Figures 2-15 and 2-16.

[2] The wire was duplexed, then quadraplexed, and to-day multiplexed in such a way that dozens of messages may speed along the same wire at the same time in both directions.[**?**]

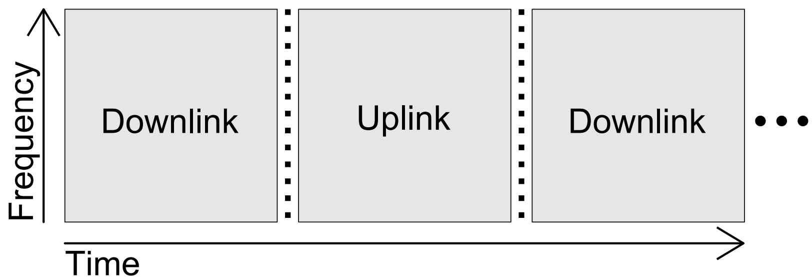

Figure 2-15: Time Division Duplexing (TDD). Transmissions in DL and UL directions are segregated by an interleaved temporal partitioning.

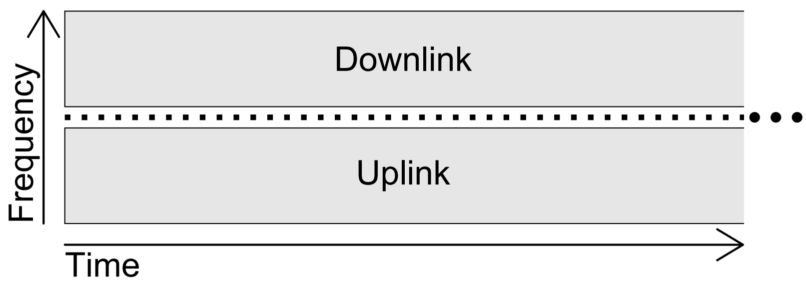

Figure 2-16: Frequency Division Duplexing (FDD). Transmissions in DL and UL directions are segregated by a frequency partition.

Multiple Access

Multiplexing is a a *physical layer* technique whereby a single shared channel is split into multiple logical channels. A single user may use multiplexing to send control and data frames along the same channel. The multiplexing offered by the physical layer is used at the MAC layer to provide a *channel access method* for users. This is called multiple access.

When TDM (Time Division Multiplexing) is used for multiple access it is called Time Division Multiple Access (TDMA), and when FDM is used for multiple access it is called and Frequency Division Multiple Access. These are the simplest multi-access approaches and are illustrated respectively in Figures 2-17 and 2-18.

OFDMA (Orthogonal FDMA) is a more recently adopted technology that allows subcarriers to be packed closer together by adopting orthonormal Fourier coefficients [32]. Strictly speaking, OFDMA is an FDM + TDM multiplexing technique which enables arbitrary frequency-time segregation of users, as illustrated in Figure 2-19.

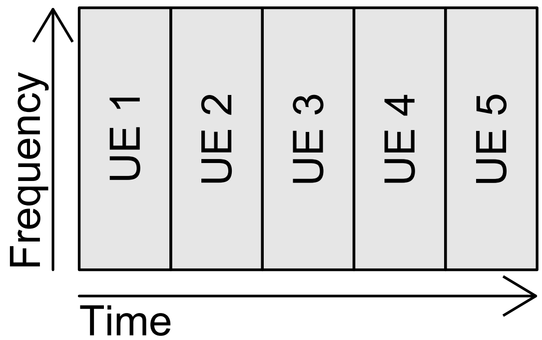

Figure 2-17: Time Division Multiple Access (TDMA). Different UEs are segregated in resource space by a temporal partitioning into *timeslots*.

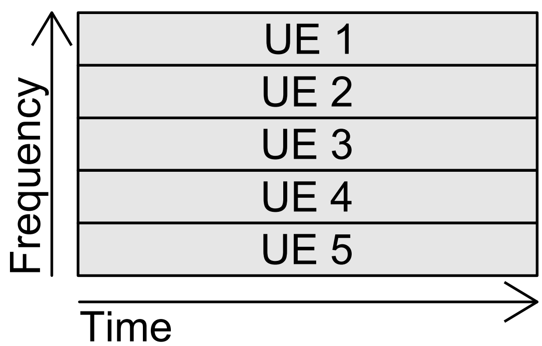

Figure 2-18: Frequency Division Multiple Access (FDMA). Different UEs are segregated in resource space by a frequency partitioning into *frequency channels*.

.

The same channel can also be shared among multiple users without FDM or TDM methods by using Code Division Multiplexing (CDM). The multiple access method is called CDMA (Code Division Multiple Access). In CDMA, each user is assigned a different pseudo-random code called a chip with which their data is bitwise multiplied, the process is called spread spectrum communications since the effect is to increase the bandwidth of the transmitted signal [33]. The codes are chosen so as to be near-orthogonal so that only bitwise multiplication with the correct code will retrieve the source signal.

In the idealised case, given two orthogonal code vectors $\mathbf{a}$ and $\mathbf{b}$ such that $\mathbf{a} \cdot \mathbf{b} = 0$, a binary sequence $\mathbf{x}$ is transmitted as $\mathbf{a} \cdot \mathbf{x}$. Decoding with the correct spreading code $\mathbf{a}$ will result in $\mathbf{a} \cdot \mathbf{a} \cdot \mathbf{x} = 1 \cdot \mathbf{x} = \mathbf{x}$ whereas decoding with any other orthogonal code, such as $\mathbf{b}$ will result in $\mathbf{b} \cdot \mathbf{a} \cdot \mathbf{x} = 0 \cdot \mathbf{x} = 0$. Note that in practice the code vector is always shorter

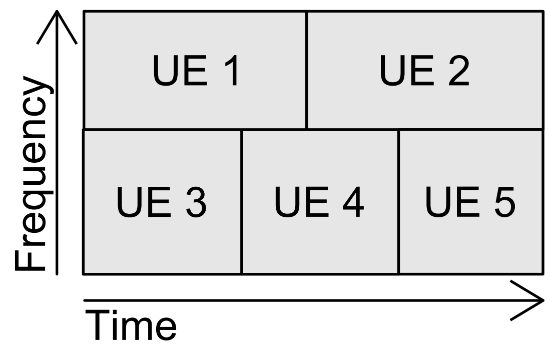

Figure 2-19: Different UEs are segregated in the resource space by blocks which span finite regions of frequency and time called *resource blocks*.

than the information vector.

The concept of CDMA is illustrated in Figure 2-20:

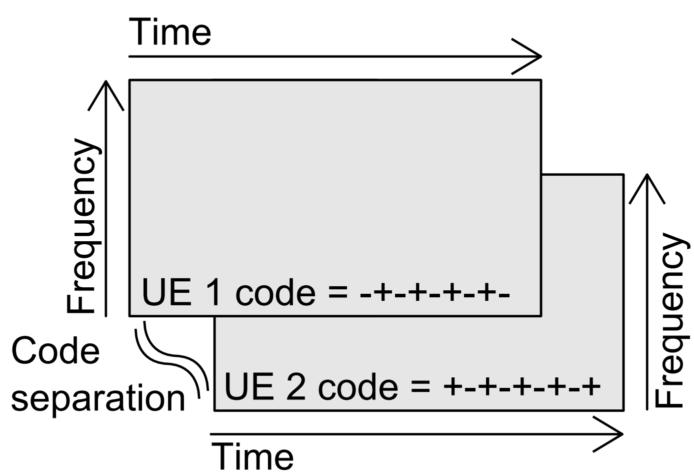

Figure 2-20: Code Division Multiple Access (CDMA). Different UEs are segregated *on the same frequency-time resource* by using different spreading codes.

Finally, the a resource can be shared by spatially separating the users. This is called spatial multiplexing. In the most obvious case UEs are registered to different cells, which provides a kind of multiplexing of resources at the system level. It turns out however, that spatial multiplexing can be applied at much smaller scales by deploying multiple receive and/or transmit antennas.

Note that spatial multiplexing can only be exploited in conditions where the different spatial paths are uncorrelated, and thus is not typically suited to line of sight conditions. Systems may employ mechanisms to try to detect if the antenna paths are uncorrelated.

Note that the minimum spatial separation of antennas is approximately half the wavelength of the transmission frequency, although in theory there is no limit. Spatial multiplexing is illustrated in Figure 2-21

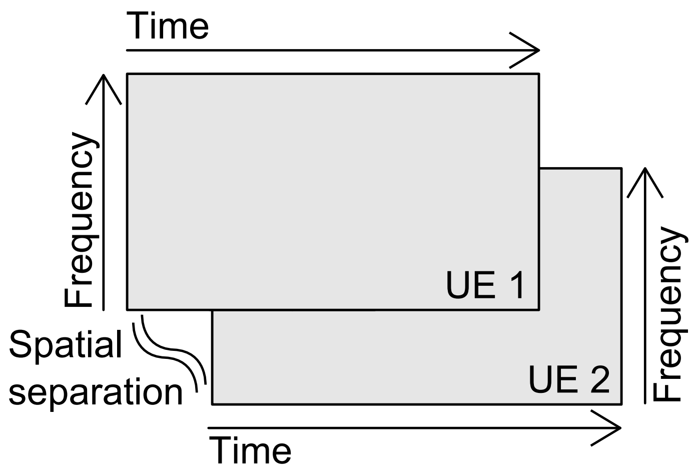

Figure 2-21: Spatial multiplexing (SM). Different UEs are segregated *on the same frequency-time resource* by using spatially separated antennas.

It is worth noting here that multiple antennas can be used in manners other than for spatial multiplexing. For example [34]:

1. STBC (Space-Time Block Codes). In this case, the multiple spatial channels are treated as a single logical channel having a greater number of resources, by using encoding schemes that exploit the offered diversity by spanning multiple spatial planes. The simplest schemes simply duplicate the data across the planes, but much more sophisticated schemes exist that can yield vast improvements in SINR. See for example [35, 36].

2. Multilayer transmission. In this case, encoding blocks do not span multiple spatial planes, but rather the additional resources are used directly to increase the data rate n-fold according to the number of additional antenna pairs used.

3. Beamforming. This refers to a preprocessing of antenna transmit parameters to compensate for predicted channel distortions. The effect is likened to a transmitter with a directable mainlobe.

All these schemes typically employ sophisticated feedback mechanisms in order to determine which to use in a given circumstance. These techniques are out of scope of this dissertation.

Modulation and Coding

Given a packet of data to transmit, the physical layer (PHY) process, will proceed along lines approximate to the following (WiMAX PHY shown) [37]:

$$Randomisation \rightarrow FECencoding \rightarrow Interleaving \rightarrow Modulation$$

Randomisation is a precursor step applied to the data to remove long repetitive sequences of 0s or 1s, affording some statistical protection against errors.

FEC (Forward Error Correction) encoding is the process of adding redundant information to data before transmission in order to increase the likelihood that it can successfully be decoded by the receiver. There exist many sophisticated techniques appearing as Reed-Solomon Codes, Convolutional Codes, Convolutional Turbo Codes, Block Turbo Coding, and Low Density Parity Check amongst others. Excellent information on the topic of coding can be found in [38]. A more basic form of forward error correction is simply to introduce repetition of the data to be transmitted followed by some majority vote or more sophisticated combining at the receiver.

The final step is Modulation. Modulation refers to the encoding of a digital signal in a form suitable for analogue transmission. The usual practice is to encode each symbol as a unique combination of phase and amplitude modulation. This is succinctly expressed by means of a constellation diagram, so called because the various combinations of phase and amplitude are shown as a constellation of points on the complex plane. Figure 2-22 shows constellation diagrams for BPSK, QPSK, and QAM-16.

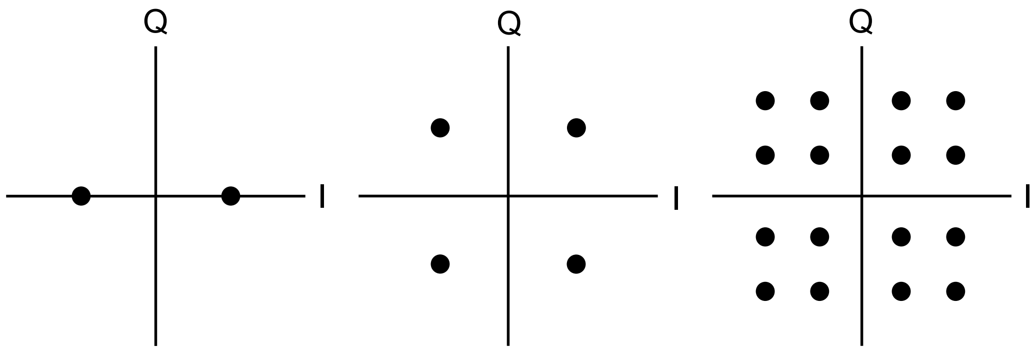

Figure 2-22: Three modulation schemes, from left to right: BPSK, QPSK, and QAM-16.

The first modulation scheme shown is BPSK (Binary Phase Shift Keying), it is called so because it uses two phase shifts to key or index two binary values. The second modulation scheme is QPSK (Quadrature Phase Shift Keying) which represents four different values using four different phase shifts. The final modulation scheme shown is QAM-16 (Quadrature Amplitude Modulation 16) which represents sixteen different values by using sixteen unique phase and amplitude shifts.

It is important to be familiar with modulation, if only to recognise it when it appears, because different technologies support different types of modulation. Note that it is also possible to perform modulation by varying the frequency of the carrier signal too.

At the physical layer, in OFDMA systems, blocks of data are transmitted using a given FEC code and Modulation scheme. Together these choices are referred to as the MCS (Modulation and Coding Scheme). The code part is sometimes referred to as the rate, since by introducing redundancy, the overall data rate is reduced. If a block is transmitted with 16-QAM rate 1, this means that no FEC was used, if the FEC adds 100% more data in the form of redundancy, then the rate will be 0.5 and so on.

The astute reader may be wondering why the number of points on a constellation diagram cannot be increased indefinitely. In theory they can, but in practice, the presence of noise blurs the points into each other, resulting in erroneous reception of the message.

Thus the amount of noise that the channel is subject to determines the highest order MCS scheme that can be supported. In good conditions, perhaps a rate 1 64-QAM can be used, whereas in bad conditions, perhaps it is necessary to use some low rate BPSK with additional repetition coding.

This is the motivation and reason why wireless transmission technologies support multiple MCS schemes. It is the job of the scheduler, for each transmitted data block, to choose the most appropriate MCS scheme given the expected channel conditions. This also then determines the data rate, for a given quantity of resource, that a UE will experience. If a fixed data rate has to be maintained for a particular UE, in the face of deteriorating channel conditions, the only solution is to increase the resources allocated to that UE in proportion to the reduction in data rate implied by the use of ever decreasing code rates and lower order MCS schemes.

How can the scheduler at the BS know the expected channel conditions for each UE? It does not, primarily because the channel conditions are constantly changing, but secondarily because it is not positioned where the UE is positioned so has a different experience.

There are two ways which a BS can estimate the channel conditions of a UE:

1. Assume reciprocal channel conditions and use its own channel estimates

   This approach is only acceptable when the following conditions hold:

   (a) The system is TDD and so the transmission frequencies used by the BS are the same as those used by the UE.

   (b) The path that the signal takes between BS and UE will be approximately the same in the opposite direction from UE to BS.

   (c) Neither the UE nor the BS are being subject to significant levels of interference from some transmitting device, that the other is not also being subject to.

   (d) The UE is stationary or is travelling slowly enough that channel conditions are not changing faster than scheduling decisions are being made.

   If all these conditions hold. Then the channel conditions experienced by the BS will be approximately the same as those experienced by the UE, and the BS can use its estimates of channel conditions to make appropriate MCS choices for the UE.

2. Accept periodic reports from the UE regarding channel conditions

   In this approach, the UE periodically sends reports to the BS about its channel conditions, and the BS uses these to make appropriate MCS choices. Note that this approach is not restricted by the same set of conditions as the former and thus is applicable to a wider range of situations, including FDD deployments. The disadvantage, is that sending reports consumes UL resources which could otherwise be employed to send UE data.

A final point that needs remarking on, is to describe exactly what information is referred to to make MCS decisions by the scheduler. What is the channel estimation, numerically, and how is it obtained?

Since received power is frequency dependent, an estimation should focus on those frequencies of interest at which data will be transmitted, and consequently for which MCS choices must be made.

MCS choices are made based on noise levels and hence the usual metric employed is signal to (noise + interference) ratio or SINR. In order to estimate SINR a receiver needs to receive a signal and then estimate how much noise it has been subjected to. In order to do this it needs the original reference signal to compare with. But how can a receiver know what the transmitted signal should be in advance of receiving it? What would be the point of transmission if it did? Of course it does not know in general but this problem is overcome by the introduction of pilot symbols.

Pilot symbols are data symbols which are known in advance of reception, as they occur on specific frequencies at specific points in time. They are programmed into the receiving device. The receiver can then compare the pilot symbol with the received symbol and produce estimates of the phase, amplitude, and noise distortion introduced by the channel. This channel estimation information is not only used as a basis for MCS selection, but is also used in the decoding process to correct channel distortion.

2.1.6 Scheduling

Scheduling is the term used to describe the process of mediating access to a shared channel. Given a set of resources, scheduling is concerned with allocating resources to UEs on a frame by frame basis. There is a large body of work in the literature providing many advanced and efficient scheduling algorithms. (See for example [39, 40, 41, 42, 43, 44, 45] for some good examples of scheduling under LTE). In this subsection we will first examine common scheduling schemes that operate in the time domain only, and show how these schemes can be modified to operate in both time and frequency domains simultaneously.

Time Domain Scheduling

For the sake of making the concept of scheduling less abstract, consider the LTE system. In LTE a scheduling decision can be made once per millisecond (as an LTE frame has a duration of one millisecond). For a cell bandwidth of 10Mhz, 50 discrete resources called Virtual Resource Blocks (VRBs) are available. It is the job of the scheduler to decide to which UEs these virtual resource blocks are assigned.

In time domain scheduling this decision amounts to choosing the UE with the highest priority, according to some scoring metric. This UE is then assigned all 50 of the available resource blocks. Alternatively, perhaps the top 5 highest priority UEs are assigned a randomly selected 1/5th of the available resource blocks each. Regardless of how many UEs are selected for scheduling, the point is that time domain scheduling decisions do not try to match up *specific* resources with UEs: the question asked is *to whom* resources will be assigned, not *which* resources each UE shall receive.

For the sake of simplicity, assume that just one UE is given resources at each timestep. As explained, the approach taken is to compute a “score” for each UE at each timestep, and to give the resources to the UE) who scores the highest. The score should reflect the mobile operator’s subscriber policy, it is not standardised, and somewhat arbitrary. Typical factors that may be considered are:

- How much time has elapsed since the UE last received data.
- How much data the UE has already received.
- The UEs estimate of the channel condition.
- The service class of the UE.

In this section an intuitive view of three common scheduling strategies will be provided. The strategies are: round-robin, greedy, and proportional fair. In the next section, these will be extended to the frequency domain. Using the approach taken in [46] schedulers can be described using a priority function. Mathematically: given N UEs, at each discrete time t, resource k is allocated to user $i*$ as follows:

$$i* = \arg\max_{1 \leq i \leq N} f_i(t)$$

where $f_i(t)$ is the score of user i at time t. The most commonly encountered scheduling approaches are Round-Robin, Greedy, and Proportional fair.

Round-robin scheduling is a very simple scheduling strategy: UEs simply take it in turns to receive resources. The priority function is:

$$f_i(t) = (i \mod t) + 1 \tag{2.9}$$

The concept is illustrated in Figure 2-23. The figure shows 9 successive frames $t0...t8$. And the letters A,B, and C indicate the assignments made. As can be seen, the assignments rotate around round-robin style, hence the name.

Round robin is inherently fair in terms of the number of resources it offers to the set of scheduled UEs, but usually results in low cell throughput relative to other schedulers because UEs in poor radio conditions are treated equally to UEs in good radio conditions. Whereas UEs with good radio conditions can actually receive a higher data rate than UEs with bad radio conditions. Whether or not a UE has good or bad radio conditions, is information usually known to it, because this information is reported to the basestation so that the basestation can make sensible MCS selections.

A UE can estimate its own channel conditions by: measuring received signal strength; measuring the degree of distortion of known reference symbols (called pilots) transmitted by the basestation; and keeping track of how many received packets have been decoded successfully. Channel quality measurements are given the name CQI (Channel Quality Indicator). In LTE, the reported CQI is actually an MCS recommendation to the basestation,

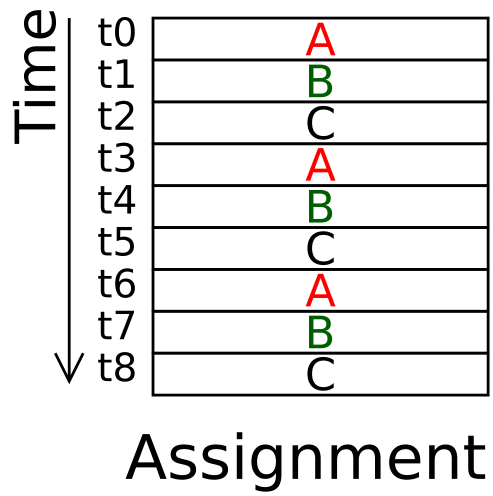

Figure 2-23: In round-robin scheduling UEs A,B, and C take it in turns to receive resources.

and thus is assumed to proxy the SINR.

The so called "Greedy" scheduler lies at the extreme end of the spectrum of schedulers that exploit channel information. In Greedy scheduling, the user with the current highest instantaneous channel condition is scheduled, the priority function is as follows:

$$f_i(t) = cqi_i(t) \tag{2.10}$$

Where $cqi_i(t)$ is the CQI of user i at time t. Figure 2-24 illustrates the concept. The figure shows for three UEs A, B, and C what their CQI score is as a function of time. This aspect is shown on the left. On the right, the corresponding assignments are shown. As can be seen, the resource is assigned to whichever user has the highest score.

The greedy scheduler certainly exploits channel knowledge but can result in the starvation of UEs with continually bad radio conditions.

In proportional fair scheduling the aim is to trade-off the exploitation of channel knowledge whilst protecting against the starvation of users. This is achieved by trading off the average throughput of the user with their instantaneous throughput, the priority function is as

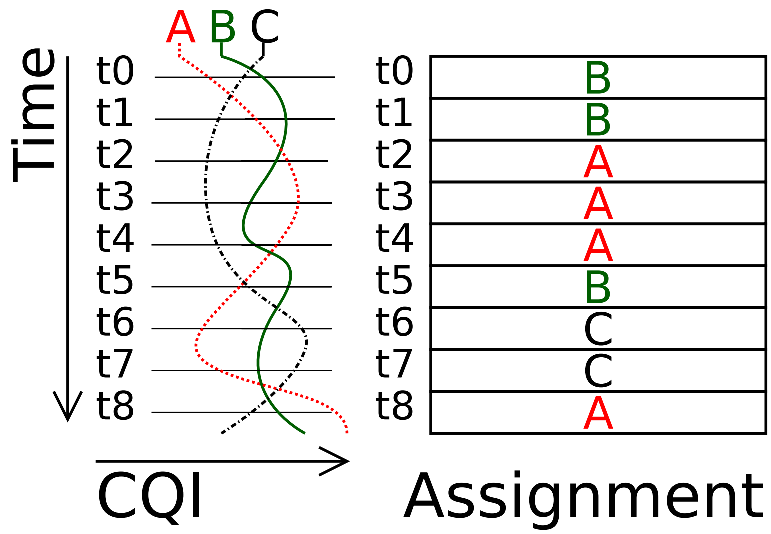

Figure 2-24: In greedy scheduling the each resource is assigned to the UE having the highest Channel Quality Index (CQI).

follows:

$$f_i(t) = \frac{cqi_i(t)}{cqi_i(average)} \tag{2.11}$$

Where $cqi_i(t)$ is the CQI of user i at time t, and $cqi_i(average)$ is the historical average of CQI for user i. The purpose of this approach is to try to take advantage of peaks in channel conditions, but only when they are large relative to the average for a given UE. A common alternative is to use mean throughput as the denominator value, in order to trade this off with instantaneous throughput. A finer control of the tradeoff can be obtained through the use of an adjustable exponent applied to either the numerator or denominator.

Frequency Domain Scheduling

Frequency domain scheduling is an augmentation of time domain scheduling, wherein at each time step, not only are high priority users selected to receive resources, but those resources are distributed selectively. The motivation for this is that the frequency-power spectrum of a received signal is not flat, meaning that some frequencies are received with higher power than others due to channel characteristics. Suitably equipped UEs can therefore estimate the channel quality of each frequency region (resource band) and report a vector of CQI measures to the basestation.

Many of the time domain scheduling algorithms are simply adapted by evaluating the priority function not only for each user but for the combination of each user and each resource band. Formally, given N UEs, frequency resource k is allocated to user $i*$ as follows:

$$i* = \arg\max_{1 \leq i \leq N} f_{i,k}(t)$$

where $f_{i,k}(t)$ is the score given to user i on frequency resource k at time t. Under this framework, a frequency domain variant of the Greedy scheduler described in the last section, can be constructed using the following priority function:

$$f_{i,k}(t) = cqi_{i,k}(t) \tag{2.12}$$

Where $cqi_{ik}(t)$ is the CQI for user i on subcarrier group k at time instant t. For each frequency resource, this scheduler picks the user who has the highest reported CQI at time t. This scheduler is commonly known as "Maximum C/I (Carrier to interference) scheduler". It is evident, under the assumption of perfect channel knowledge that a scheduler employing this priority function cannot be exceeded in throughput.

The proportional fair scheduler described above can be extended to the frequency domain via the following priority function:

$$f_{i,k}(t) = \frac{cqi_{i,k}(t)}{cqi_{i,k}(average)} \tag{2.13}$$

Where $cqi_{i,k}(t)$ is the instantaneous CQI of user i on frequency resource k at time t and $cqi_{i,k}(average)$ is the average CQI of user i on frequency resource k. As in the time-domain case, a common variant is to use the average throughput of user i as the denominator value.

It should be intuitively obvious by now that schedulers can be constructed to trade off channel conditions and fairness, in many different ways.

2.1.7 A Historical Overview

From a handful of core technological approaches, a large number of esoteric technology variants have evolved and been practically deployed across the world. The literature contains experiments and recommendations based on both the generic technological approaches, and their various progeny.

It is important to understand this distinction since the similarities between an approach and its relatives determine the transferability of derived conclusions.

If experimental conclusions are based on a generic approach then it needs to be ascertained whether these conclusions hold for the restricted and variant progeny of that approach. And on the contrary, if experimental conclusions are based on a restricted form of some generic approach, it needs to be ascertained whether the conclusions hold for restricted forms.

Furthermore, some results and optimisation approaches may generalise across generic approaches, and thus it is important to differentiate between approaches and understand their differences.

Thus, whilst it would be superfluous to document the history of mobile phone technology in its entirety, it is worthwhile to provide a brief synopsis for the purposes described above. In outline, three generations of technology have existed and the current focus is on the fourth generation. These generations are as follows:

- First generation (1G)

  1G systems are differentiated from all subsequent technological generations in having utilised analog channels. All modern cellular technology uses digitally modulated signals.

1G systems were circuit switched, which means call resources were setup in advance of a call and maintained throughout it. The only adjustment to the circuit resources occurred during handover to another cell.

Resources in this case were channels: a pair of transmission frequencies. One of the pair was used for communication from phone to basestation (uplink) and the other was used for communication from basestation to phone (downlink).

There was a small degree of scope for optimisation, the candidates were: positioning of basestations, basestation parameters, initial cell selection, handover, and transmit powers.

A myriad of incompatible 1G systems existed, and international roaming was not possible.

Almost all practically deployed 1G systems have been obsoleted and replaced with 2G systems, and they are scarcely mentioned in contemporary research. Thus no further shall be said on the matter.

- Second generation (2G)

  When speaking of second generation technology, or 2G, this invariably means GSM (Global System for Mobile communications).

  GSM is uses both TDMA and FDMA technology. FDMA is used to segregate the frequency bandwidth available into a set of 200khz wide frequency channels. Each frequency channel is broken down using TDMA into frames which last 4.615ms and contain 8 timeslots each having a length of $557\mu s$. For a phonecall, a UE is allocated one timeslot within a TDMA frame within one channel. GSM is generally deployed in FDD configuration [8].

  GSM can transmit packet based data in addition to phonecalls via the extension GPRS (General Packet Radio System) which is also referred to as 2.5G. This is enabled by allowing access to timeslots on a request basis as opposed to a circuit switched call basis. This packet based access has been further improved by the development of GSM EDGE (Enhanced Data rates for Global Evolution) in the year 2000 [47]. GSM EDGE uses 9 MCS schemes: GMSK with 4 different code rates, and 8-PSK with 5 different code rates [9]. GSM EDGE Evolution, the latest incarnation of the technology, extends this to 12 MCS schemes, introducing QAM-16 and QAM-32, in addition to allowing dual channel data transmissions [47].

GSM is widely deployed in every global commercial market and is still heavily used.

- Third generation (3G)

  Global 3G development is driven by an international government organisation called the ITU (International Telecommunication Union) [48]. The project under which contemporary 3G technologies fall is called IMT-2000 (International Mobile Telecommunications 2000) [49]. IMT-2000 provides spectrum recommendations, and data rate requirements, which have mostly been adopted.

  Actual IMT-2000 compliant technical specifications have been developed by a number of different external bodies. The most important of these are 3GPP [13] (3rd Generation Partnership Project) UMTS and 3GPP2 [50] (3rd Generation Partnership Project 2).

  3GPP is a partnership between European (ETSI [51]), Japanese (ARIB [52] and TTC [53]), American (ATIS [54]), Korean (TTA [55]), and Chinese (CCSA [56]) standards bodies. It has developed an ITU-2000 W-CDMA standard called UTRA and this standard is what is meant when the term 3G is used in Europe.

  3GPP2 is a partnership made up of ARIB, TTC, TTA, CCSA and a North American partner (TIA [57]). 3GPP2 was borne out of a technology war and has developed a 3G standard incompatible with UTRA called cdma2000.

  Focusing on UTRA since this is the technology widely deployed in Europe, the technology is based on Wideband CDMA (W-CDMA), which just means the system is based on CDMA and transmits over a relatively wide bandwidth. The technology is often referred to as just UMTS. When speaking of the network aspects, the term UTRAN (UTRA Network) is often used.

  UMTS comes in both TDD and FDD variants. In UMTS FDD, each frame lasts 10ms and is split into 15 slots, giving a slot length of $0.666ms$.

- Fourth generation (4G)

  For marketing reasons, 4G systems are not usually referred to as such, but appear instead under the names of their technologies. The most important 4G systems are WiMAX and LTE, which are both based on OFDM.

  WiMAX is standardised under IEEE 802.16e [58], and a major certification effort is underway at the WiMAX forum [59]. The largest player in its development has been

Intel [60]. Clearwire [61] and Sprint Nextel [62] have been promising its widescale U.S deployment since 2007 but have yet to realise this vision.

LTE Release 8 [63] is standardised by 3GPP group, and is the current 4G technology of choice for Vodafone and other European operators. LTE differentiates itself from WiMAX in its frame structure and the way that channels are encoded.

The work in Chapter 3 is based on WiMAX, whereas the work in Chapter 4 and Chapter 5 is based on LTE Release 8.

2.1.8 Simulation overview

In this section an overview of the simulation platform used by the experiments described in Chapter 3 and Chapter 5 is provided. This information is given because it is supposed that it will provide useful context for the description of the research work which follows. It is not intended to describe precisely the simulator or the parameters used in simulations. All relevant parameters and system simulation constraints are described explicitly wherever experiments that utilise the system simulator occur. Note that the work in Chapter 4 follows a Monte-Carlo approach and uses the simulator only to generate pathloss and shadowing information, it does not use any of the dynamic features described below.

Simulation initiation

The first task performed by the simulator is to establish the network environment. There are two major tasks:

1. Establish the hexagonal layout, and the position and directionality of antennas.
2. Precompute pathloss and shadowing information for the simulated area.

Figure 2-25 illustrates these ideas by depicting a hexagonal network with a single tier of sites and 21 cells. As can be seen, the position and directionality of antennas has been established. Underlying the hexagonal deployment, a grid can be seen, this grid indicates the resolution to which pathloss and shadow fading are computed.

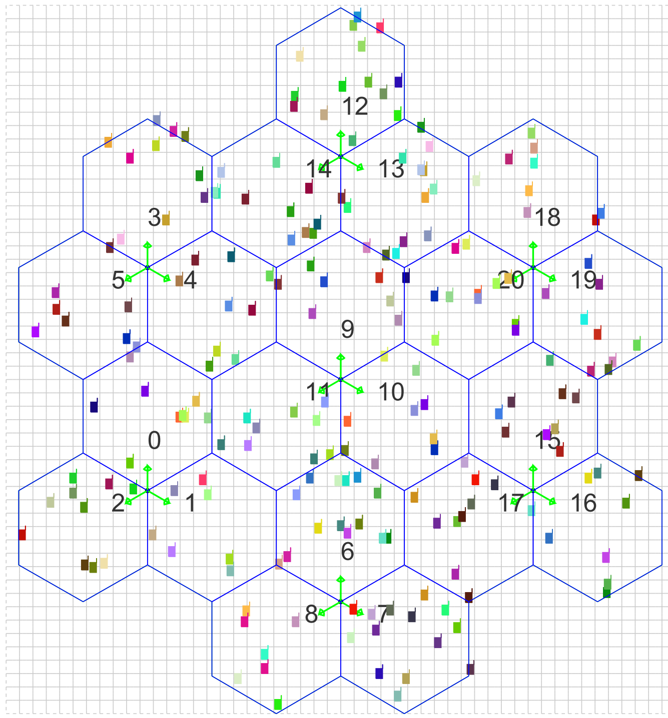

Figure 2-25: A single tier hexagonal deployment of sites, as represented by the system simulator.

For each grid position, pathloss plus shadow fading *from each cell* in the simulation is computed and stored. This means that the resolution of this information is determined by the resolution of the grid. In the figure, a $30m^2$ grid resolution is used for illustration, but it is typical to use a resolution of between $5m^2$ and $10m^2$ for satisfactory accuracy.

Once the grid has been established, the users are dropped on to it. Users are illustrated in the figure by mobile phone shaped objects. Users are usually positioned randomly in such a manner as to establish a given number of users per cell, or per unit area.

Each user is associated to the basestation with the strongest signal, according to the precomputed grid information. The following attributes are constructed for each user:

1. Mobility model

   Each user has a dedicated mobility controller, which is called upon periodically to update their position. It characterises the user's velocity and direction change probabilities. In the simulations presented in this dissertation, users do not explicitly move about the simulation area, but rather their "mobility" is captured via the statistics of fast fading channels which are parameterised to reflect their supposed velocity.

2. Fast fading channels

   For each possible pairing of basestation and user in the simulated area, a fast fading channel is constructed. It is assumed therefore that, for the purposes of the model, the fast fading profile applied to transmissions from basestation x to user k is independent of the fast fading profile applied to transmissions from basestation y to user k. Note that a user *decodes* transmissions from only one basestation at a time, but receives *interference* from all transmitting basestations, which is why a fast fading channel is needed for each basestation-user pair. The fast fading information characterises the impact of the user's velocity on their radio state, and is used to scale the precomputed pathloss data. It is updated periodically.

3. Data channels

   Each user has a dedicated data buffer and associated traffic model. This implies the construction of several auxiliary structures including: MAC PDUs, RLC PDUs, HARQ processes if necessary, and transport block templates.

4. Measurement trackers

   Each UE makes periodic measurements of various network aspects, such as the RSRP (Reference Signal Received Power) of each sector, and the measured CQI on each subband for the serving sector. Data structures to keep track of this information must therefore be established. In all simulations in this thesis, exponentially decaying averaging windows are used with an exponent of α, such that new data x is incorporated into the moving average a_t at time t as follows: $a_t = (1 - \alpha)a_{t-1} + \alpha x$. In this work α assumes a value of 0.5.

Simulation Execution

The system simulator is a discrete time simulator. The temporal resolution is set according to the length of the TTI (Transmission Time Interval) of the radio technology of interest.

A TTI is the minimal permissible time between one transmission and the next. In the case of LTE examined in Chapter 5, this interval is $1ms$.

The system simulation proceeds discretely for each TTI until the desired number of TTIs have been simulated. The following actions are performed within each simulated TTI:

Update user attributes

The first task is to update the user attributes, so that user state information is consistent with the supposed events that have occurred between the previous TTI and the present one. The following attributes are the most important:

1. Fast fading

   Since exponentially decaying moving averages are used to track CQI, historical information is forgotten quickly. However since some historical information is present, this means that CQI is not tracked perfectly, and thus the speed of the fast fading channel has a bearing on performance aspects that depend on being able to track it. Updating the fast fading every TTI is thus an important step.

2. Measured attributes

   Rapidly changing measures such as per-subband CQI are updated for each user.

3. Position and handover

   The user position is updated according to its associated mobility attributes, and the handover status is examined, if the necessary criteria are met, a handover is performed.

4. Traffic

   The user's traffic model is updated. To give an example, if a bursty traffic model is used, then updating the model might result in a new packet appearing in that user's packet buffer.

Scheduling, transmission, and reception

Once the user attributes are updated, the next step is to schedule data to users and perform transmission and subsequent reception. Scheduling is performed for all users that are active and have data to send. The actual algorithm used varies. A basic introduction to scheduling was provided in Section 2.1.6. The result of scheduling is the production of transport blocks that will be sent over the air interface using a given MCS.

After scheduling, transmission and reception is simulated. Packets are assumed to arrive near instantaneously, and thus transmission and reception happen within the same TTI. When a transport block is received, the mean SINR on that block is computed. The signal is determined by the shadowing at the user's current position and the fast fading channel between it and its serving sector. Interference is computed using the shadow fading and fast fading of all transmissions that do not original from the serving sector. Noise is a fixed quantity.

Given the SINR of the transport block, the packet error rate can be obtained using a map produced by link-level simulations. Given the packet error rate, a random number is generated to decide whether the packet is received in error or not. The threshold is set so that the mean result is consistent with the packet error rate obtained from the lookup curve.

Successful reception of a packet propagates the received data up the user's protocol stack, such that their packet buffer reflects the current state. Simulation statistics for dropped packets, user throughputs, and so on are then updated.

2.2 Frequency Selective Scheduling in 802.16e Mobile WiMAX

This part of the literature review pertains to the work presented in Chapter 3. That work examines the impact of changing velocity on the performance of frequency selective scheduling in 802.16e Mobile WiMAX. Section 2.2.1 provides a technical overview of 802.16e and introduces the terms and features used and exploited in Chapter 3.

Section 2.2.3 reviews existing work that examines the relationship between user velocity and frequency selective scheduling, and provides the motivation for the work carried out in Chapter 3.

2.2.1 WiMAX 802.16e Primer

Frame structure

The mobile WiMAX 802.16e TDD frame structure [37, 64] is shown in Figure 2-26.

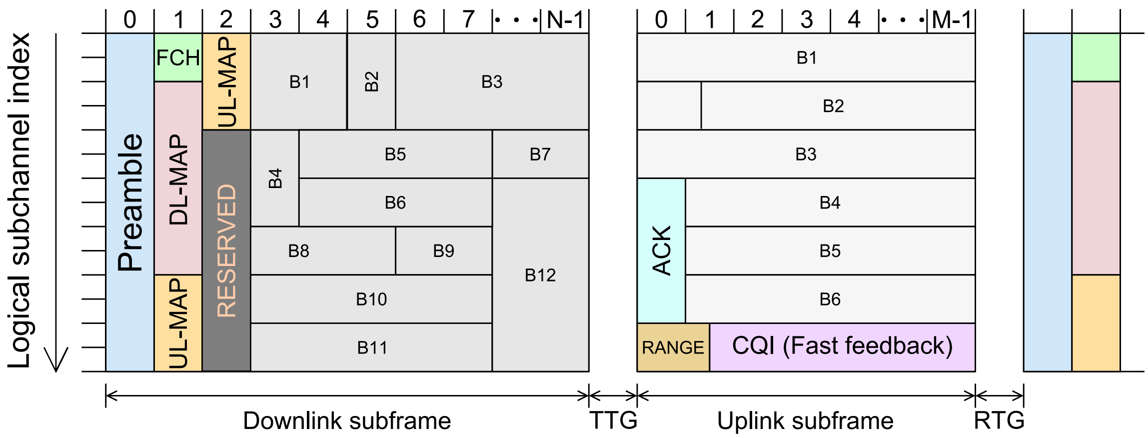

Figure 2-26: 802.16e Mobile WiMAX frame structure

A frame is split into downlink and uplink subframes, separated by a small Transmit/receive Transmission Gap (TTG). Each frame is separated by a small Receive/transmit Transmission Gap (RTG). The gaps are considered part of the frame, whose total duration is 5ms.

Preamble occupies the first OFDM symbol in the frame and spans all subcarriers. Preamble contains known reference signals and is used for synchronization, as part of RSSI (Received Signal Strength Indicator) and CINR (Carrier to Interference + Noise Ratio) estimation, and cell identification.

FCH (Frame Control Header) occupies the first two logical subchannels, and spans one OFDMA symbol. It contains information used to decode the DL-MAP and the UL-MAP.

The DL and UL maps specify the positions of bursts and the users allocated to those bursts in the downlink and uplink frames respectively. In DL, a burst is a 2D resource with its dimensions specified over subchannels and OFDMA symbols. In UL a burst is a 1D resource specified by a start and stop symbol; the burst occupies all symbols in-between.

The size of bursts can vary within a sub-frame and is a MAC layer decision implemented at the level of the scheduler. The smallest allocatable burst equals one slot which is 48 data subcarriers. Every burst is therefore an integral multiple of this. Each burst is allocated using a single MCS.

There is a tradeoff between burst size and signalling overhead in the DL map, because with smaller bursts, a greater variance in MCS can be achieved, to the possible benefit of better throughput, but at the expense of requiring more signalling information in the DL-MAP to describe the burst structure. On the contrary having fewer bursts reduces the signalling overhead but also reduces the MCS granularity to which users can be allocated, thereby possibly reducing throughput. The point at which throughput minus overhead is maximum is the optimal compromise.

Users should be grouped into MCS categories and packed into bursts appropriately. A user identifies his MAC PDU (Protocol Data Unit) within a multi-user burst by his unique MAC CID. It is typical to use one MAC PDU per burst [65].

CQI (Channel Quality Index) is a field which is used to feed-back measured subchannel SINR to the basestation so that this information may be used to scheduling advantage by setting transmission powers, selecting MCS schemes, and subcarriers appropriately. RANGE is used to feed back information for initial allocation of transmit powers and their subsequent adjustment.

2.2.2 Zone structure

Introduction

In WiMAX, the minimum resource unit is termed a slot. A slot always contains 48 data subcarriers. The WiMAX standard defines several schemes by which slots are constructed, where they differ in the way that subcarriers are drawn from the subcarrier space, the number of OFDMA symbols that a slot spans in time, and the number and position of pilots. Different schemes have been designed to serve different purposes for both the uplink (UL) and downlink (DL). These are now described in detail:

Different slot structures and associated subchannel frameworks are called zones. Figure 2-27 shows all the zones defined in the WiMAX standard. WiMAX profiles further qualify which zones are mandatory to implement for a given technology, and which are optional.

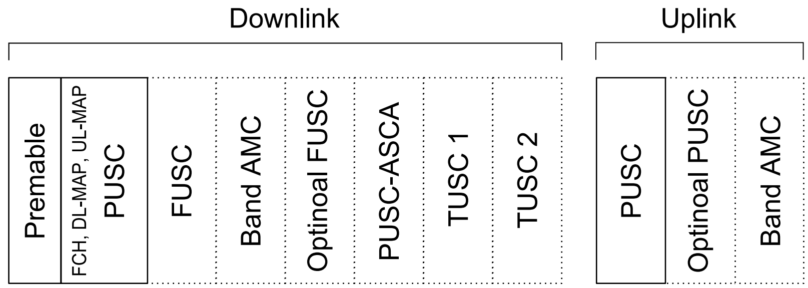

Figure 2-27: WiMAX permutation zones.

FUSC

In FUSC (Full Use of Subcarriers), a slot equals one subchannel over one OFDMA symbol. The slot contains 48 data subcarriers, they are selected by a Reed Solomon permutation of the physically indexed subcarriers. This is illustrated in Figure 2-28

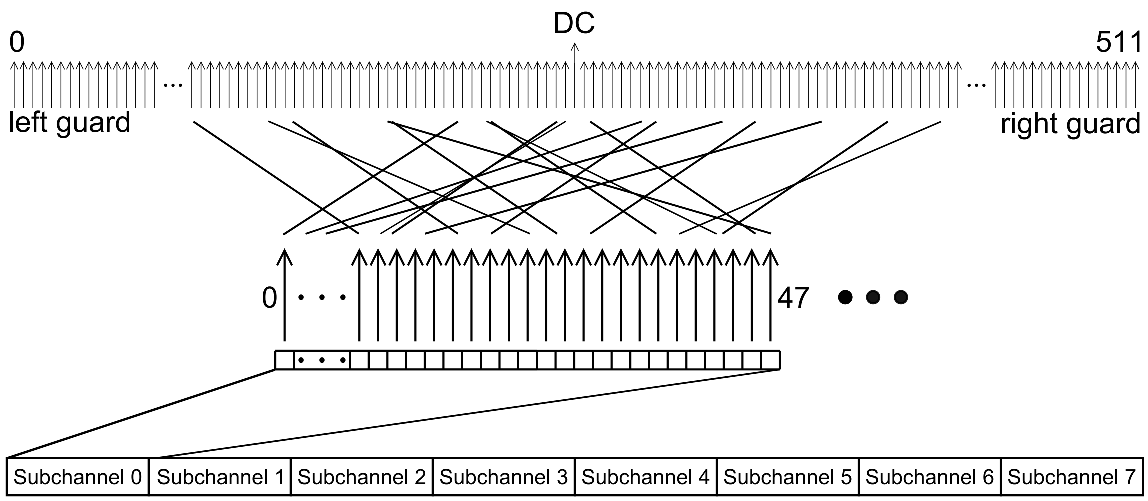

Figure 2-28: FUSC permutation. Consecutive subchannels contain 48 data subcarriers each, the data subcarriers are selected by a pseudo-random Reed Solomon permutation. Pilot subcarriers are not shown.

Band AMC

In Band AMC (Adaptive Modulation and Coding) permutation, subcarriers within each slot are contiguous. This means that UEs can more accurately estimate the state of a slot, and by providing feedback to the BS, each slot can theoretically be allocated to the UE with the best estimation, thereby allowing more accurate MCS selection and a consequent improved throughput.

AMC slots are formed from groups of contiguous subcarriers called bins. A bin contains 9 contiguous subcarriers, the middle subcarrier is a pilot, and the four subcarriers either side are data subcarriers. Six bins are combined to form a slot. There are four ways to form a bin, as illustrated in Figure 2-29.

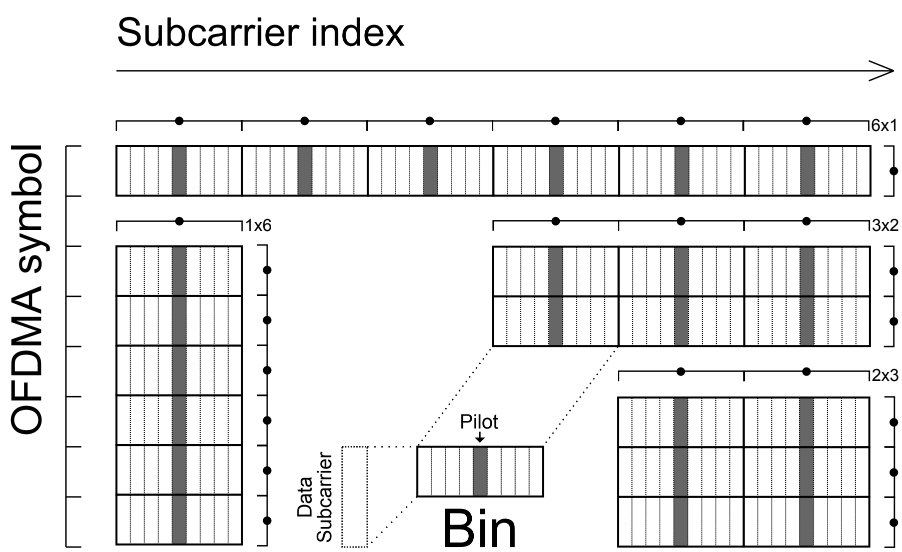

Figure 2-29: Band AMC subchannel types

AMC 2x3 is the only AMC slot type supported in version 1, revision 1.4.0 of the Mobile WiMAX Profile [66].

PUSC

In PUSC (Partial Use of Subcarriers), subchannels are constructed in a more complicated manner so that subchannels end up in different groups which can then be allocated to different sectors to obtain different reuse factors.

The process is illustrated in Figure 2-31 (Page 76), for 512 subcarriers. Here are the steps in detail:

1. The bandwidth is divided into subcarriers, the subcarriers form:
   (a) A left guard region.
   (b) Several consecutive physical *clusters* of 14 contiguous subcarriers each.
   (c) A central DC carrier.
   (d) More consecutive physical *clusters* of 14 contiguous subcarriers each.
   (e) A right guard region. The size of the guard regions, and number of clusters, depends on the FFT size and is specified in the standard.
2. The physical clusters are numbered consecutively from left to right (PC_0 ... PC_N).
3. The physical clusters (PC_0 ... PC_N) are permuted according to a scheme specified in the standard to give renumbered clusters RC_0 ... RC_N.
   (a) The renumbered clusters are allocated in consecutive blocks to six groups. Depending on the FFT size, some groups may contain no clusters.
4. Within each group, subchannels are formed from *constructed* pairs of clusters. The clusters are logically constructed as follows:
   (a) The pilots in logical cluster n are are taken from renumbered cluster n such that the pilot structure remains unchanged. Note however that the positions of the pilots is different for odd and even symbols and is shown in Figure 2-30.
   (b) The data subcarriers for logical cluster n are drawn from all the renumbered clusters within the group according to a permutation function specified in the standard that depends on n.

Thus there are two levels of permutation, the first assigns physical clusters to groups according to the renumbering scheme, and the second permutation allocates subcarriers within each group to subchannels created from pairs of logical clusters.

At the end of this process different groups are allocated to different sectors. And at least one group must be allocated to each sector.

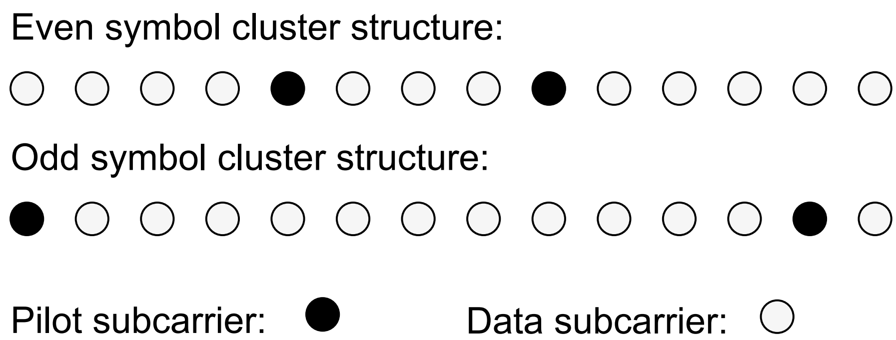

Figure 2-30: PUSC cluster structure.

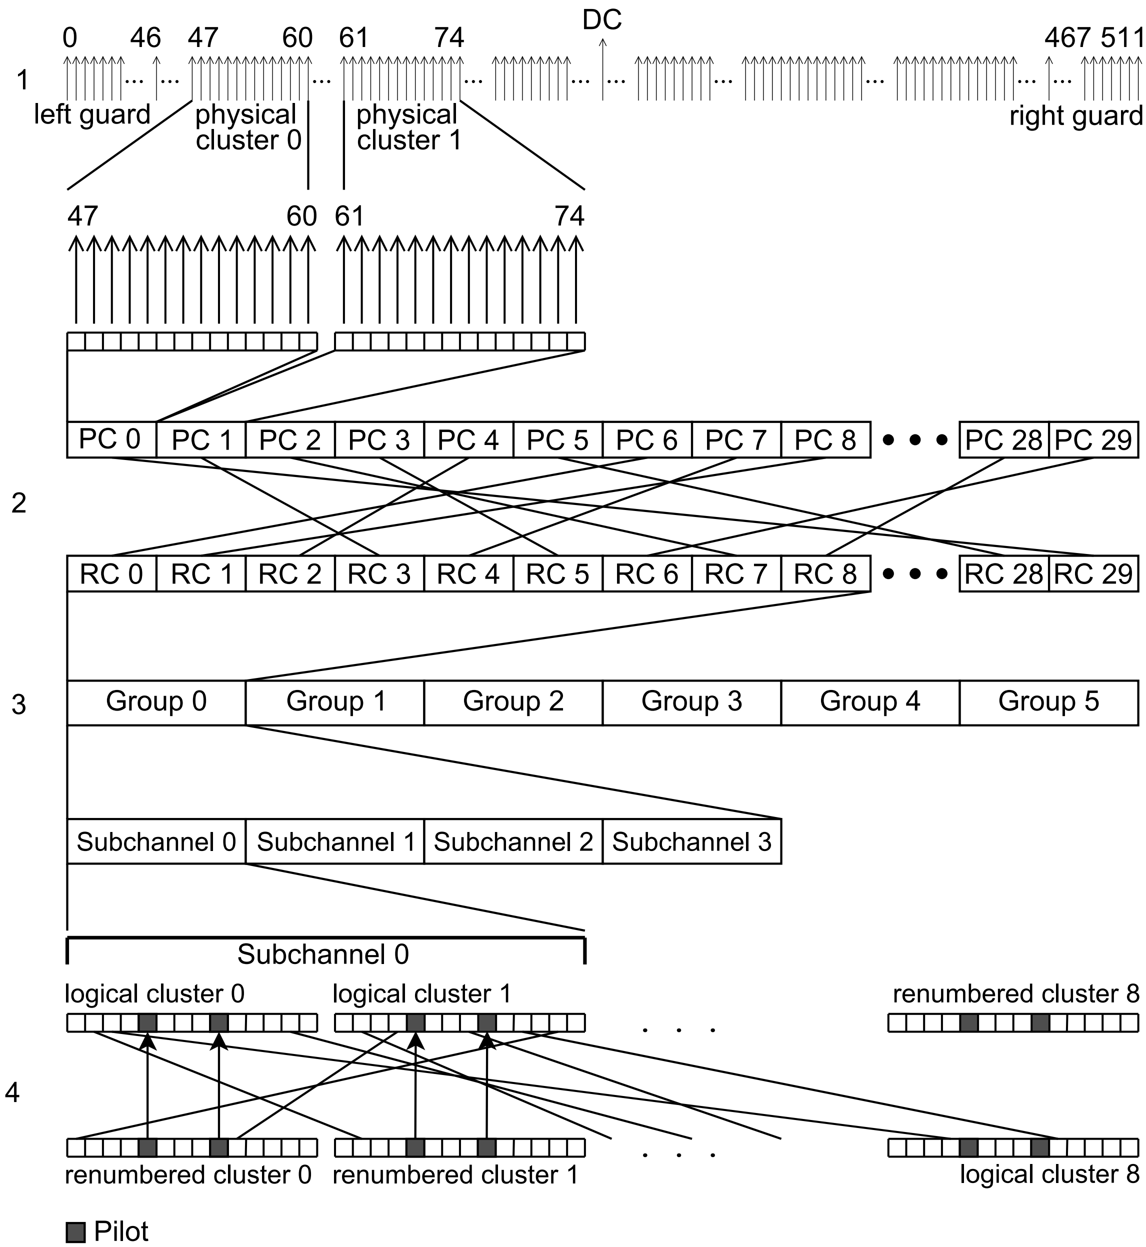

Figure 2-31: PUSC permutation.

2.2.3 The Impact of Mobility on Frequency Selective Scheduling

Given the background on WiMAX provided above, it should be clear that WiMAX offers a number of "zones" within a frame, that have different properties. The work in Chapter 4 compares Band AMC mode with PUSC mode, at different MS velocities.

At the time the work in Chapter 4 was executed, three related works occurred in the literature. In [67], based on a theoretical analysis, values on which to base switching from frequency selective scheduling to frequency spreading, are provided in the form of channel time-autocorrelations, for a metric of outage probability, and for a generic OFDMA system. It is therefore difficult to see how the results can be mapped to practical recommendations for 802.16e.

In the second [65], simulation of 802.16e is executed for both the link and system level. Average sector throughput is compared for frequency selective scheduling and frequency diverse scheduling for various dynamic and static admission thresholds. Frequency selective scheduling is shown to have a greater average sector throughput than frequency diverse scheduling at both 3km/h (by $\approx$18%) and 30km/h (by $\approx$5%), although in the latter case the outage probability is higher by 3%. The use of admission thresholds are found to trade-off total throughput and outage probability.

The results suggest that frequency selective scheduling may be relevant for vehicular traffic classes at 30km/h, although it is difficult to tell if there are any significant differences with such small margins. No other velocity classes were tested so a general picture cannot be formed, it could be that shortly after 3km/h the performance of frequency selective scheduling tails off quickly and stays relatively constant upto the 30km/h point, or it could be that a general degradation is observed. This motivates the need to study more velocity points.

The third work [68] comes to different conclusions than the second; a 50% gain in cell throughput for frequency selective scheduling relative to frequency diverse scheduling is demonstrated at a MS velocity of 3km/h, and yet at 8km/h Frequency Diverse Scheduling is demonstrated to have a 5% improvement in cell throughput relative to Frequency Selective Scheduling. It is not clear whether this discrepancy is due to error in either one of the last two studies, or from the confounding effects of the scheduler employed by the last, or just natural variance of the techniques under differing conditions.

Given this state of affairs, the work in Chapter 3 was executed to focus purely on the effect of velocity on decoder performance, without complicating the issue with additional factors. A greater range of velocities was studied than either [65] or [68] so as to give a better picture of how the relative performance of the two schemes changes with velocity.

Given that the assumptions in Chapter 3 are quite optimistic, alignment with the results of [68] rather than those in [65] seems more probable in practice. The work supports the recommendation in [22], that frequency selective scheduling should be used by default in LTE, especially since LTE has both quicker CQI feedback cycle, and a finer resource granularity than WiMAX.

2.3 Interference Coordination in LTE

The purpose of this chapter is to review preceding work on interference mitigation, and to discuss it in relation to the novel work of this dissertation. The chapter is split into three sections. Section 2.3.1 introduces the problem, Section 2.3.2 reviews static solutions to the problem, and Section 2.3.3 reviews dynamic solutions to the problem. The reviews of static and dynamic solutions provide the necessary context and motivation for the work presented in Chapter 4 and Chapter 5, respectively.

2.3.1 The Interference Problem

Interference is a problem in cellular networks because large geographical areas must be covered using finite quantities of spectrum from transmitters that have finite and imperfect spatial footprints with soft boundaries. Interference occurs at the interfaces between these footprints unless signals are transmitted orthogonally in frequency, time, or codespace.

The notion of interference mitigation has been around since the dawn of wireless communications. In Tesla's 1903 patent [69] a method to address the problem of "disturbances emanating from other sources, as well as the signals or messages being received by instruments for which they are not intended" is advanced. The patent describes a rotating wheel which produces "a series of impulses or disturbances differing from each other in character and order of succession". This is nothing other than the now familiar concept of frequency hopping: in order to avoid prolonged interference on a given transmission frequency, the transmission frequency is periodically changed so as to average out the effect of randomised interferers.

In addition to frequency hopping, the first modern narrowband wireless systems such as GSM, sought to avoid interference through mutually exclusive assignment of frequency resources among potential interferers. In GSM, an SINR of approximately 9dB is required in typical urban conditions to deliver full rate speech [7]. This high network entry requirement means that only a low level of interference can be tolerated, with the result that GSM systems were and still are deployed in frequency reuse configurations [8]. An example of frequency reuse is shown in Figure 2-32.

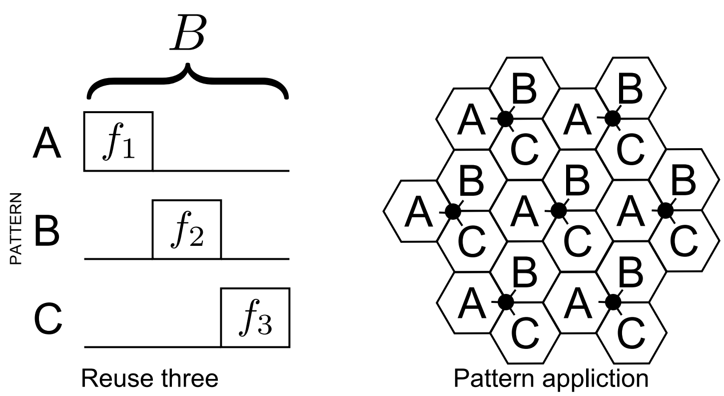

Figure 2-32: Frequency reuse with a reuse factor of three. The three patterns on the left are applied to a network of cells as on the right such that adjacent sectors use different frequencies.

The figure shows the available bandwidth B on the left split into three parts to define three patterns: A, B, and C. These patterns are applied to a network of cells as in the right of the figure, such that adjacent cells use different transmission frequencies: f_1, f_2, and f_3. This concept is called frequency reuse, and in this case the reuse factor is three. Notice that the pattern tessellates. Frequency reuse improves interference but has the obvious drawback of reducing the bandwidth available in each cell in proportion to the reuse factor.

The high SINRs required by GSM for normal operation, *necessitated* frequency reuse. LTE, on the other hand, can operate at SINRs as low as -6.5dB [70] through appropriate selection of MCSs. This lead to the suggestion during the standardisation of LTE that it could be deployed in a reuse one configuration, where all of the bandwidth B is available to use in each cell. And Vodafone is now deploying LTE equipment using a reuse one configuration in Germany (Throughout 2011).

However, even though LTE *can* operate at low SINRs (relative to GSM), it is not desirable to do so. A low SINR must be met with an MCS that employs a high degree of redundancy; reducing the bits per symbol carried. A tradeoff therefore exists: if adjacent cells use the same transmission frequencies, then although more bandwidth will be available per cell, interference will be high, and so more resources will be required for a given transmission. On the other hand, if adjacent cells use different transmission frequencies, then less bandwidth will be available per cell, but interference will be reduced, and thus less resources will be required for a given transmission. This tradeoff between bandwidth and interference is well known. It appears under the name "compensation effect" in [71].

More generally, interference coordination techniques may trade some other commodity for reduced interference. For example, a reduction in transmission power in one cell may alleviate the strain on the interfered in another, but it comes at the cost of weakened signals in the donor cell. These kind of tradeoffs are at the heart of investigations into interference avoidance, and the net effect of them, taken in the context of operator policies, determines the efficacy of a given technique.

In this section we examine various approaches to interference mitigation found in the literature. The analysis is split into two main sections: static interference coordination and dynamic interference coordination. The former concerns techniques which do not adapt to network changes whereas the latter concerns techniques which do.

Note, interference coordination goes by many names in the literature, for example: interference mitigation, interference compensation, interference avoidance, and interference reduction. Some common mnemonics are: ICO for Interference Coordination, and ICIC for Inter-Cell Interference Coordination. In this work, these terms are used interchangeably.

2.3.2 Review of Static Solutions

The simplest interference coordination techniques are static in nature, meaning that, a certain permanent configuration is established within a network that it is assumed will provide the necessary improvement in SINR required for satisfactory operation. This configuration is kept unchanged for the lifespan of the network. Static approaches have the advantage of requiring no inter-cell communication, although they can be costly to deploy because neighbour relations must be known in advance and the cells configured accordingly. Semi static schemes usually employ similar techniques, but change the parameters on demand or infrequently at timescales much greater than the coherence time of the radio channel, for example: daily, weekly, or monthly.

On the Differences between Downlink and Uplink

When it comes to static and semi static interference mitigation, the downlink and uplink directions differ slightly. The differences come about because by definition semi-static and static solutions attempt to leverage coordination benefits that are independent of fast fading. This means that the SINRs they seek to change are dominated by the pathloss and shadowing components of transmission.

The relative positions of interfering transmitters are thus important, which leads to a difference between the downlink and uplink problems. In the uplink direction, the transmitters change position whilst the receivers remain fixed, whilst in the downlink direction the situation is reversed.

Figure 2-33 shows two adjacent cells with sector antennas X and Y. Two UEs are depicted in each cell as A,B and C,D respectively.

Consider the uplink transmissions $A \rightarrow X$, and $D \rightarrow Y$. Since the uplink direction uses power control, and since A and D are both near to their serving sectors, they will both transmit at a low power. Given that A and D are also spaced far apart from each other, their uplink transmissions are unlikely to interfere significantly.

On the other hand, consider the uplink transmissions $B \rightarrow X$ and $C \rightarrow Y$. Since B and C are both far from their serving sectors, they will transmit at a high power. And given that they are also very close together, their uplink transmissions *are* likely to interfere significantly.

There is a benefit therefore in ensuring that B and C are not scheduled using the same resources in the uplink. This benefit, however, is not present in the downlink direction.

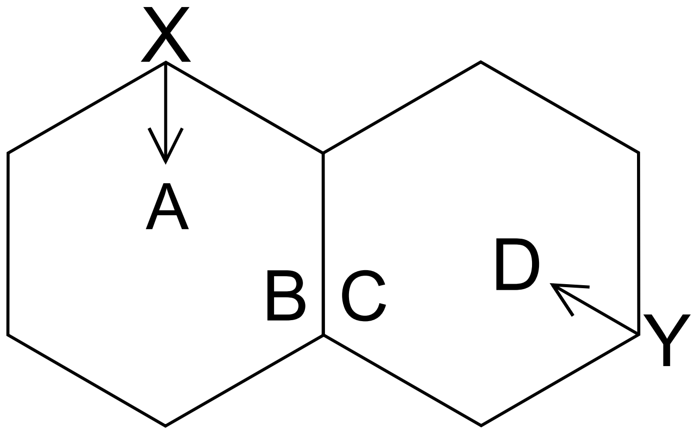

Figure 2-33: A scenario to illustrate the differences between pathloss dominated uplink and downlink interference mitigation. See text for details.

This is because the positions of the transmitters X and Y are always the same regardless of to which pair of UEs they are transmitting.

The only exception is when the resources R_{xb} used by X for B and R_{yc} by Y for C are used on a mutually exclusive basis. That is, when R_{xb} is used only by X, and R_{yc} is used only by Y. This is because *any* use of R_{xb} by Y will affect user B independently of the position of the target UE of Y. In this way, the interference problems for downlink and uplink differ (in the absence of fast fading).

Soft and Partial Reuse Schemes

Figure 2-34 illustrates, drawing inspiration from [72], two popular static ICIC schemes known as partial reuse and soft reuse.

In partial frequency reuse, the available bandwidth is split into two regions: B_1 and B_2. Region B_1 contains only one frequency band f_1 which is available to use in all cells. Region B_2 is split into three frequency bands f_2, f_3, and f_4 and a reuse three pattern is applied to the effect that the overall network uses reuse one on B_1 and reuse three on B_2. The size of the different regions B_1 and B_2 as well as the number of frequency bands B_2 is split into, can in principle be changed.

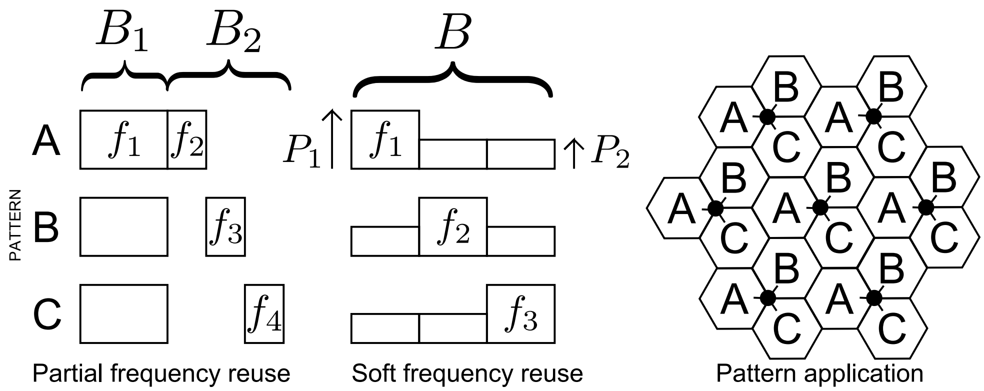

Figure 2-34: An illustration of two popular static inter-cell interference schemes: partial frequency reuse and soft frequency reuse. The bandwidth usage patterns on the left are applied to the network as on the right to create tessellating configurations.

In soft frequency reuse, the whole bandwidth B is available to use in each cell but two different transmission powers P_1 and P_2 are applied to the three frequency bands f_1, f_2, and f_3. As the figure illustrates, in patterns A,B, and C the frequency bands f_1, f_2, and f_3 respectively are transmitted at the higher power P_1 and the remaining frequency bands are transmitted at the lower power P_2. The idea is that the boosted parts should receive less interference since they are both transmitted at a higher power and are coordinated with adjacent cells. The power factors P_1 and P_2 can be changed and in principle so can the number of frequency bands.

The origin of these techniques can be traced back to a GSM technique, proposed by Nokia Telecommunications, called underlay-overlay. The idea was to split the available bandwidth into two frequency bands having different reuse factors: one lower than the other. Intelligent variants of the idea such as [73] suggest that UEs should be placed into the lower reuse band only when their C/I exceeds some threshold, and switched back to the high reuse band when their C/I drops accordingly. In this way the UE is supposed to be placed in a band that is suitable to its current interference level. The approach is very similar to partial reuse, the contemporary incarnation of which appeared first in [74].

Soft frequency reuse was proposed in its current form by Huawei in [75] during early discussions by the 3GPP during the development of LTE. These discussions concerned the avoidance of interference through higher frequency reuse factors. The industrial motivation for such techniques received significant attention, and similar proposals appeared from Ericsson in [76], LG Electronics in [77], and Alcatel in [78]. Semi-static variants of soft reuse were proposed first in [79]. Thus, it was principally from these discussions that the soft reuse concept emerged.

Simulation of Soft and Partial Reuse

In [75] it is claimed for the downlink direction that soft reuse offers "Improved spectrum efficiency and cell-edge data rate" relative to reuse one. The work however does not define cell-edge data rate nor provide any quantitative measurements concerning it. The presented results show only that orthogonal spectrum offers higher SINRs than shared spectrum, and that introducing frequency reuse reduces the total cell throughput due to a reduction in bandwidth. These obvious points provide no information as to the efficacy of soft reuse in practical situations.

In [80] Alcatel present simulation results for their scheme in the uplink direction. The results are a straight mapping between SIR and throughput, and are examined for a UE travelling between two cells. The results show only that SIR is improved under frequency reuse. The same conclusion can be drawn for partial reuse in the uplink as presented in [81]. The latter goes so far as to say "A straight implementation of frequency reuse 1 ... is unlikely to fulfil the challenging cell edge throughput requirements for the EUTRA uplink" yet no relevant throughput values are presented.

In [82] both partial and soft reuse are examined and compared against reuse one in the downlink direction. A number of different scenarios are examined: sometimes an increase in cell-edge throughput is shown in trade off for reduced total throughput relative to reuse one, however sometimes both cell-edge and total throughput are shown to reduce relative to reuse one. Partial reuse fairs much better and the results claim that straightforward tradeoff between total throughput and cell-edge, relative to reuse one, can be obtained.

The mean cell throughput under soft reuse for the downlink is examined in [83]. It is claimed that 5th percentile throughput can be improved relative to reuse one in trade off for a reduction in total cell throughput by applying soft reuse. However, the work only examines the mean cell throughput, and provides no insight into the behaviour of soft reuse under more realistic schedulers.

In [84] a load adaptive partial reuse scheme is compared with Huawei's soft reuse scheme (which they call Ericsson's scheme), Siemen's partial reuse scheme, and Alcatel's variant of soft reuse in the downlink direction. The authors conclude that their results "proof (sic) that the scheme is effective in improving the whole network's performance". But the figures in the work show a great degree of variety between the schemes tested for different performance metrics, such that it is not clear which is the overall winner. Furthermore, the schemes are not compared against reuse one, so no insight into the relative performance over the baseline case can be obtained.

In [72] a mixture of reuse three and reuse one is examined for the downlink direction, whereby users are scheduled in a round robin fashion. At each TTI the scheduled user is given either the whole bandwidth or one third of the bandwidth at three times the power. This decision is based on an SINR threshold, such that UEs with SINRs below the threshold are given the boosted one third, and UEs with SINRs above, the whole bandwidth. Bizarrely, the scheduling actions of adjacent cells are made completely independently, rendering the claim that the approach is a mixture of reuse one and three completely meaningless. The authors conclude that reuse one obtains better 5th percentile and total cell throughputs than either the mixture or reuse three alone.

In [85] it is claimed that reuse three improves SINR compared to reuse one, at the cost of reducing the aggregate throughput to 75-80% of the reuse one rate. However, no data or references are provided to substantiate this claim.

In [86] soft reuse in TDD LTE is examined, whereby users are assigned to cell edge or cell centre bandwidth regions based on a distance metric. It is claimed that no improvements in 5th percentile throughput can be obtained using soft reuse, relative to reuse one, for full buffer traffic. It is also claimed that improvements in total throughput can be obtained using soft reuse, relative to reuse one, for bursty traffic; in this case a relatively low bandwidth VOIP (Voice Over IP) model.

Both soft reuse and partial reuse are examined in [87] in the downlink for static deployments without fast fading and compared to reuse one and reuse three. Those UEs with the worst geometry factor are given the boosted resource for the soft and partial reuse cases. UEs are divided into cell-centre and cell-edge classes so that they receive the same power per VRB. The work claims that reuse one obtains higher total throughput than any of the compared schemes, but suffers a proportional drop in 5th percentile throughput.

Discussion

No clear conclusions can be derived from the examined work as to the efficacy of either soft frequency reuse or partial frequency reuse, relative to reuse one in downlink or uplink directions.

The presented results are not only unclear and sometimes unsubstantiated [85, 81], but are also conflicting. For example: both [82] and [88] claim that partial reuse, relative to reuse one, obtains improvements in throughput at the 5th percentile point, yet [89] concludes that “the basic partition-reuse scheme studied was not capable of improving the rate at the 5% CDF point”.

As another example, [90] claims that soft reuse provides gains in both cell-edge *and* total throughput when compared with reuse one yet [91] concludes that “With the expected link performance no improvement can be found with static downlink reuse schemes.”. On the other hand [82] differentiates itself by demonstrating losses for soft reuse in both cell-edge *and* total throughput for some scenarios, whereas [86] presents a mixture claiming that soft reuse can provide gains for bursty traffic but not for full buffer scenarios.

The discussions in [89] and [91] go a little way to *explain* their own results, and remark on the interaction between SINR changes and bandwidth to this end. They do not however explicitly enumerate the factors which conspire to produce the disparate results seen. Neither do any significant attempts to offer general explanatory analysis appear elsewhere.

For these reasons, it was considered necessary to begin the investigation into interference mitigation with the simulation of static soft reuse schemes. The first contribution of this dissertation, within the field of interference mitigation, is to offer a precise enumeration of the factors involved in determining the outcome of statically applied soft reuse. With a little imagination it isn’t difficult to see how the conclusions also apply to partial reuse or more general static schemes. This analysis appears in Chapter 4.

Another problem with the body of work discussed above is that none of the presented results are statistically sound, in the sense that none of the quantitative data provided is qualified by means of tests for statistical significance. All of the presented results, with the exception of [82], draw conclusions based on single-shot, so called $n = 1$ simulations. Or at least, that is the impression obtained given the methodology descriptions. This casts the validity of all these results under a shadow of doubt.

The work presented in Chapter 4 addresses this concern by qualifying all presented results with appropriate tests for statistical significance, and in accord with expert advice [92], transparency with regard to associated p values.

Another manner in which the work of this dissertation departs from its predecessors, is in providing results for non-hexagonal deployments. The question of how non-hexagonal deployments may react differently to static interference mitigation is thus raised. The work in Chapter 4 shows that gains obtained using tessellated patterns in hexagons, tend to deteriorate under heterogeneous deployments. The deterioration is not absolute however, but is enough to raise doubt in any hexagonal results that show only *marginal* gains.

A frequently seen assumption in the work discussed in this section is that boosted or orthogonal resource regions should be given to UEs at the cell edge. For example, in [87] soft reuse is examined, and the boosted resource is given preferentially to UEs at the cell edge. The motivation for this is usually that the cell-edge UEs suffer more from interference than the cell-centre UEs and thus deserve the boosted resource more than the cell-centre UEs. But this doesn't take into account that *all* UEs typically benefit from using a boosted resource, even the cell-centre UEs. It is interesting to ask therefore whether it is indeed optimal to give boosted resources preferentially to the cell-edge UEs, or whether it might be better sometimes to give the boosted resources to the cell-centre UEs. The work in Chapter 4 treats this question seriously, by challenging the assumption, and shows for a fixed rate service class, that it is better to give the boosted resource to UEs at the cell-centre.

Given this illustrative example, Chapter 4 goes on to show more generally that the conditions under which a given scheme is optimal, depends precisely on the scheduling strategy adopted, such that the successful deployment of a static scheme would require consistent operator policy and user distributions. It is due to this revelation that the work moves onto dynamic solutions.

2.3.3 Review of Dynamic Solutions

In Section 2.3.2 it was revealed that UE position has a greater impact on interference in the uplink direction than in the downlink direction. This was explained to be because uplink transmitters (i.e UEs) can have varying absolute positions whereas downlink transmitters (i.e eNBs) have fixed absolute positions. This assertion assumes that the pathloss components of interfering signals are the dominant components. Static solutions consequently tend to investigate fast-fading independent gains which manifest only in the mean over relatively long periods of time and across multiple UEs.

Introducing fast fading, means that downlink and uplink interference is no longer dominated by pathloss, since the variation of SINR due to fast fading is much greater. Since LTE will almost certainly employ frequency selective scheduling to track this variation for low velocity UEs, and will obtain substantial gains from doing so [22]. It is therefore necessary to ask in which ways interference can be mitigated in a dynamic setting where fast fading, as well as UE distribution changes, can have a significant impact on performance.

Dynamic schemes, like static schemes, mitigate interference through two principle means: (i) prudent use of resource orthogonality when it is possible and beneficial and (ii) intelligent selection of shared resources when it is not. Techniques need to show gains above and beyond those already present due to scheduler exploitation of frequency and user diversity to be worthwhile.

The work of [71] provides a useful insight in to the types of gains that may be possible. Uplink simulations are performed which focus on the tradeoff between using a greater number of shared resources against using a lesser number of orthogonal resources, where the latter benefit from higher SINRs that the former. The term "compensation effect" is coined to describe the compensation of reduced SINR with additional bandwidth. Simulations are performed between pairs of cells in a multi-cell environment to study this tradeoff.

The first results they present, albeit obvious, are something of a portent. They claim that when bandwidth utilisation is low, and resource usage is unrestricted, then enforcing orthogonality is not usually worthwhile since there are plenty of spare resources to compensate for the reduction in SINR caused by collisions. On the other hand it is claimed that when bandwidth usage is low, but resource usage is restricted, then orthogonality can be beneficial because the UEs do not have enough freedom to compensate for the reduction in SINR caused by collisions.

Their data suggests that for $\approx 90\%$ of resource collisions between UEs in adjacent cells, UEs are better off using a greater number of shared resources than a lesser number of orthogonal resources. And although the scenario examined is somewhat simple, if there is any merit in the results, this suggests that leveraging gains from coordination of transmissions in high interference scenarios is likely to be difficult.

From the opposite side, it follows that for $\approx 10\%$ of the UEs, enforcing orthogonality at the expense of bandwidth results in a net gain. It is probable that this 10% is itself dynamic. A multitude of sophisticated techniques are discussed in this section that try and capitalise on this percentage.

Interestingly, the authors of [71] seem to overlook this latter conclusion in that they do not speculate on exploiting this percentage in full interference scenarios. I.e. in situations where all resources would otherwise simultaneously be used in adjacent cells, they overlook that it may be beneficial to selectively enforce orthogonality.

The remainder of this section discusses selected dynamic techniques which try to address the problem and concludes with a discussion.

Centralised Solutions

An obvious question that arises when considering sophisticated interference mitigation solutions is the extent to which cells should communicate and at what level of abstraction scheduling decisions should be made [93]. At one extreme, each cell acts independently of the others and uses inter-cell communication only for handover, whereas at the other extreme all cells communicate frequently with a central controller that reciprocates the arrangement by sending out periodic instructions.

In this section a overview of centralised solutions is provided through the analysis of selected examples.

Graph based approaches In [94] a centralised solution based on graph theory is presented for the downlink direction. A fully connected graph where UEs form the nodes, is weighted such that the weight between two UEs is proportional to the interference they would experience if simultaneously scheduled on the same resource, assuming that this interference can be estimated accurately. Two cell centre UEs will thus tend to have a lower weight than two cell edge UEs, since the latter suffer more from collisions than the former. A maximum k-cut of the graph is then executed, where the size of the cut is determined by the sum of the weights of the edges that are cut. The k-cuts correspond to k resource bands and assignment proceeds accordingly. This tends to schedule cell edge UEs in different bands. The approach is promising, but is compared only to a completely random scheduler so provides no information about how effective the technique is in comparison to schedulers that take user fairness and frequency selectivity into account.

Another approach that draws on graph theory is found in [95]. A "conflict graph" is established such that nodes represent UEs and edges exist between UEs in different cells to which simultaneous transmission would cause significant interference. The existence of an edge implies that the connected UEs should not be scheduled on the same resource. The graph is constructed such that strict adherence will result in all UEs meeting a specified target SIR. Resource assignment follows by repeated random Hamiltonian cycles of the cells in the network, scheduling one UE per cell visit. Round robin switching is used among UEs upon successive visits to the same cell. At each cell visit, the current high priority UE is assigned a 3x12 block of 802.16e subchannels. Once a frequency block is used, it may not be used by any other conflicting UE within that TTI (Transmission Time Interval).

The approach described in the last paragraph is compared with reuse three. Both approaches are examined with and without a beamforming implementation that assumes perfect UE tracking to a resolution of 1°. The comparison however does not seem entirely fair: in the reuse three case, the round robin scheduler samples all UEs, whereas when the conflict graph is in use, only a subset of UEs are sampled due to limitations imposed by assignments which are marked as forbidden in the conflict graph. Since lower SINR UEs tend to have more connections, this suggests that the algorithm, despite visiting UEs in random order, essentially favours higher SINR UEs because they are less likely to be blocked by conflict connections early on. Unfortunately then, the only useful conclusion is that (idealistic) beamforming can improve throughput by $\approx 50\%$ compared to no beamforming.

In [96] the authors revisit the same algorithm, and examine the impact of restricting the operating distance of conflict graph construction, i.e, whether or not far away UEs are included in conflict graph relations. The work claims that that the improvement obtained when switching from zero to one tiers of coordination is larger than the improvement obtained when switching from one to two tiers of coordination.

The paper examines a soft reuse scheme which is coupled with zero tier coordination (intrasite coordination), and claims that it out-performs zero tier coordination alone. The truth is, both schemes are shown to offer a series of tradeoff points, such that what is best is a matter of operator policy. The presented tradeoffs are not particularly useful in any case since they fail to consider frequency diversity. There is nothing to refute the proposition that the tradeoffs examined could be obtained by tuning the parameters of frequency selective scheduling under reuse one.

In fact, evidence which supports this latter position is seen in further work by same group. In [97] and [98], which appear as almost identical works, the algorithm is revisited yet again. This time, conflict graph construction distance is set to one tier only. It is claimed that this is almost as good as global coordination (a somewhat superfluous claim if we believe the previous paper [96] which makes the same claim of zero tier coordination). The algorithm is then examined for what is effectively different degrees of scheduling fairness. The overall conclusion is that by changing the balance between scheduling high SINR UEs (vertex degree is used as an inversely proportional proxy for this) and other UEs, that the tradeoff in 5th percentile and total throughput can be controlled. This is quite interesting in light of the sentence ending the previous paragraph, and resonates with the work presented in Chapter 4.

Two tier approaches In [99] the authors introduce an inter-cell interference coordination algorithm that operates at two levels: globally and locally (the authors point out that the global level could in principle be regional). At the global level, optimal assignment of channels to UEs is sought by maximisation of the aggregate throughput. The throughput equations however consider whether or not a pre-identified dominant interfering eNB is transmitting on the proposed channel or not. Assignment then proceeds by visiting the eNBs in ascending order starting with the eNB with least assigned channels. At each visit the channel and UE which improve the aggregate throughput the most are paired, taking into consideration any interference from already assigned dominant interferers. This step is executed at a so called superframe temporal granularity, which is an order of magnitude longer than that of normal scheduling decisions which are made once per frame. At each TTI, every eNB follows the prescription laid out by the global algorithm, unless the selected UE has nothing to send, in which case the channel is reassigned to another UE using a maximum C/I approach. Whilst the advanced approach is interesting, it suffers the same fate as others, in being compared only to a random scheduler. Furthermore the experiments are performed for a hexagonal grid with only 7 cells; not a very challenging interference environment.

The same global algorithm is examined and compared to a sector independent maximum C/I scheduler in [100] and demonstrates a maximum of 36% throughput gain over the latter for an average bandwidth occupancy of 30%. These gains diminish to 1% as the bandwidth occupancy is increased to 100%. It is not clear why a maximum C/I scheduler is used, since the approach is unrealistic and unfair. This makes it difficult to draw any conclusions about the efficacy of the proposed technique in practice. It also brings to light that the global part of the technique is unfair, in the sense that its global decisions are greedy. This is likely to be undesirable for a network operator in the same manner that a local maximum C/I scheduler is. In addition to the discussed work, the paper also examines a priority based interference avoidance strategy identical to that used as a benchmark in Chapter-5, where it is given the acronym TD-PF-SS. The comparison is performed against a random scheduler however so again little insight can be gained. It is worth noting that the work also presents a variant of the priority scheme that gives additional priority to low SINR UEs. The work remarks that the 5th percentile is improved compared to the forebears. It should be pointed out that this latter advancement is nothing more than a scheduling bias, a point which resonates with the comments made in Section 2.3.3.

In [101] an interference mitigation algorithm is presented that leverages inter-cell communication on a temporal scale comparable to the channel coherence time. A dual tiered algorithm is presented which operates at the sector level and at the global level. In the first stage each sector can choose to restrict two dominant interferers per UE per chunk. That is, for a given chunk the two most dominant interfering neighbours in the first two tiers may be proposed for silencing. This decision is based on the predicted gains from doing so being larger than a fixed threshold. The gains can be calculated because the authors assume that each UE can disassemble the SINR on each chunk into its neighbourly contributions by using band-specific RSRP measurements.

Once each sector has calculated restrictions for each UE, the Hungarian algorithm is used to optimise the proposed mapping of resources to UEs (the Hungarian algorithm is an efficient solution to the combinatorial optimisation problem of perfect bipartite graph matching [102]). These optimised assignments are then sent to a central controller, whose job it is to resolve conflicts in proposed transmission restrictions. Conflicts are resolved by an integer linear programming solver which maximises the joint utility. The utility measure contains a factor proportional to the inverse of the number of imposed restrictions, so that the cell with least restrictions is favoured, introducing an element of fairness. The CDF of user throughput shows gains at all percentiles relative to Hungarian scheduled reuse one and thus on paper the approach would seem to be an excellent idea.

A slight modification of the technique is compared to soft reuse, partial reuse, and reuse one in [103]. The differences are that: the global optimiser no longer considers the number of imposed restrictions (thereby removing the fairness element), different thresholds are used for deciding whether a neighbour should be restricted, and the catchment of neighbour restriction is limited to one tier. It is unsurprising therefore that the CDFs look slightly different when compared with those of the former exposition, but in general the conclusions are the same. An incremental improvement is however shown for an extension which allows restricted PRBs (Physical Resource Blocks) to be used at a lower transmit power instead of completely forbidden.

Finally in [104], the authors modify the algorithm again to incorporate a more realistic scheduler (proportional fair) and study the performance for full buffer and bursty traffic. Interestingly, the performance of reuse one is now presented as having the highest total throughput whereas in [101] and [103] the novel scheme was. That aside, under the best novel scheme presented, relative to reuse one, the total throughput drops by $\approx 1Mbps$ in exchange for a 5th percentile throughput increase of $\approx 0.35Mbps$. It is not clear whether this tradeoff could not therefore be obtained by tuning the proportional fair parameters to favour the 5th percentile in the reuse one case.

Whilst the approach presented in the three papers [101, 103, 104] looks very promising, it unfortunately suffers some significant shortcomings. The most pressing is that the algorithm assumes that band and eNB specific RSRP (Reference Signal Received Power) is periodically reported by UEs. This has to be done as frequently as band specific CQI is measured for it to be of any use, and is likely to have a similar overhead. Whilst this overhead is not insurmountable, it should certainly be taken into consideration. A graver and probably unassailable problem is that 3GPP LTE *does not support* band specific RSRP reporting. From the perspective of practical implementations this means that the approach as described is not suitable for extant deployments. Whilst reciprocity could be used in a TDD system coupled with proprietary backhaul communication to overcome this problem, the focus here is on FDD systems, because these represent Vodafone's immediate concerns.

It is also interesting to ask of the three papers, how fair the centralised aspect of the algorithm is, since while cells with the least restrictions are favoured in [101], this condition disappears in [103] and [104]. It is not clear how this affects the fairness of the algorithm, since the variance is not reported. It is possible that the mean gains in 5th percentile throughput observed come at the cost of increased inter-cell variance and reduced inter-cell fairness.

Distributed Solutions With Communication

It is briefly worth mentioning that in 3GPP LTE release 10 there are three standardised inter-cell communication mechanisms. These are as follows:

1. UL OI (Uplink Interference Overload Indication)

   This measure is defined for the uplink and is communicated by a list of entries, one for each VRB, to neighbours over the X2 interface. Each list entry is an enumerated type and the number corresponds to the degree of interference generated on the given VRB. For example, the enumeration may confer values such as (high interference, medium interference, low interference, etc). The idea is that a cell will tell its neighbours in advance that it will be causing high interference on a given VRB, probably because it is scheduling to cell-edge UEs on these resources. The neighbour cells can then avoid using these resources. This measure is defined in Section 9.2.17 of [105].

2. UL HII (Uplink High Interference Indication)

   This measure is defined for the uplink and is communicated by a binary vector, with one element for each VRB, to neighbours over the X2 interface. A value 0 in the position for a VRB indicates "low interference sensitivity" and a value 1 indicates "high interference sensitivity". The idea is that a cell can tell its neighbours which resources its cell-edge UEs are using in the uplink, so that they may avoid using the same resources for their own cell-edge UEs. This measure is defined in Section 9.2.18 of [105].

3. RNTP (Relative Narrowband TX Power)

   This measure is defined for the downlink and is communicated to neighbours over the X2 interface. It consists of binary vector with one element for each VRB, and a threshold. A value 0 indicates "Tx not exceeding RNTP threshold" and a value 1 indicates "no promise on the Tx power is given". The idea is that neighbours can know which VRBs are being transmitted at a power lower than the threshold, and can in principle increase their own TX power accordingly in a manner similar to soft reuse, or use this information in some other way, such as avoiding those resources for which no guarantee on TX power is given. This measure is defined in Section 9.2.19 of [105] and the threshold is defined in Section 5.2.1 of [106]

Note that HII and OI are actually quite similar, the only difference is that HII says "you are causing me interference on VRB X", whereas OI says "I am potentially causing you interference on VRB X". But the response of both is really the same, in that the ideal neighbour reaction is to avoid using resource X.

A scheme which takes advantage of HII, and promoted as a dynamic version of soft reuse, is proposed in [107]. It is called adaptive soft reuse and examined for the uplink direction. UEs are classified into cell edge and cell centre based on RSRP and respectively assigned resources from either the soft part of the spectrum or the reuse one part of the spectrum. It is assumed that the proposed cell-edge assignments of neighbours are available via HII messages, and based on this, the scheduling assignments are rearranged to avoid PRB collisions. Since the cell edge and cell centre resource regions are assigned exclusively to their respectively classified member UEs, a saturated cell edge region prevents further cell-edge UEs from being scheduled. To combat this, cell-edge resources in adjacent cells are used when the absence of HII messages indicates their availability. The scheme and a variant are compared against ordinary soft reuse and reuse one. Unfortunately, reuse one is examined using random resource assignment, and so no useful insight can be obtained into the performance of the proposed schemes against more realistic schedulers.

A distributed scheme is presented in [108] that assumes communication of load information between cells. Given this information, each cell decides independently how many block interleaved frequency resources it should use, relative to its neighbours. The scheme is shown to improve packet delay when compared with reuse three. It is not clear whether the benefits seen are simply due to the statistically higher resource ceiling obtained, or whether the orthogonal nature of the resource splits is the dominating factor.

A similar approach is advanced in [109] where load information is again exchanged between cells. In this case, the cell-edge load of a given cell relative to its neighbours determines its share of a common resource pool, such that cells with more cell edge users get more resources. This is compared to soft frequency reuse, and reuse one. It is claimed that the proposed scheme improves the 5th percentile throughput by between 10 and 50 kbps compared to soft frequency reuse, with lesser improvements being obtained at high loads. Soft frequency reuse obtains a 5th percentile improvement of between 50 and 300 kbps compared to reuse one, although at high load suffers a total throughput reduction of 300kbps. It isn't clear whether at high load, reuse one could obtain the same 5th percentile throughput simply by increasing the priority of the cell edge users.

Another load based interference mitigation scheme is advanced in [110]. The scheme is quite simple: each cell estimates its load based on mean UE SINR and buffer status. An overloaded cell may then request its neighbour to turn down its transmit power. If this request is granted, the overloaded cell can increase its transmit power proportionately. A version which requires no communication is also proposed, whereby underloaded cells proactively reduce their transmit power whenever possible (i.e power control). The first scheme is compared against reuse one and it is claimed that it can reduce the mean UE buffer size in sparsely populated areas.

The scheme is similar to that proposed in [111] which itself draws inspiration from [101]. In this case, the utility of each UE is monitored with respect to a QoS target. If a UE remains unsatisfied for some threshold period, it is flagged. The eNB may then propose to restrict resource usage in one dominant interferer per UE, such that the interferer will either be banned from using the resources of the flagged UE, or can only use those resources at a lower power. The results presented indicate 5th percentile throughput improvements as well as mean cell throughput improvements relative to reuse one.

It seems improbable however that the claim is true, since improving the throughput of a given UE requires restriction of resources in another cell. This is unlikely to result always in a net benefit, given that the instantaneous channel conditions in the adjacent cells will vary. Indeed, careful observation of the figures presented in the work, shows that the gains of the approach diminish with increasing load. This implies that the gains seen at lower load are due to restrictions placed on adjacent cell resources which are *not needed* in the adjacent cell.

Taking a different approach, two algorithms are presented in [112] that operate on groups of cells. The work takes the example of three cells, and in particular it examines the cells at the intersection of three adjacent sites. The concept is illustrated in Figure 2-35.

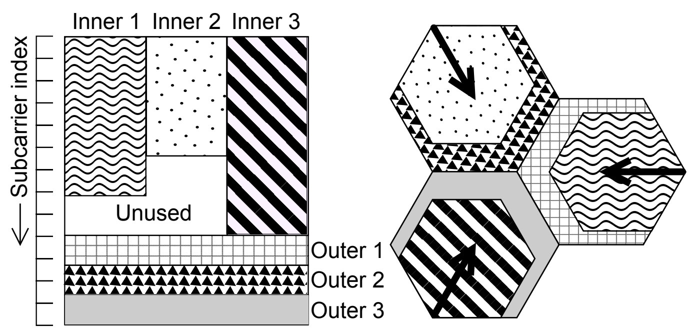

Figure 2-35: See text for details.

UEs are classified into near and far groups as a function of SINR, defining inner and outer zones for each sector. The first algorithm allocates resources to inner UEs first, starting from subcarrier index 0 in each cell, as in the left of Figure 2-35. It then allocates resources to outer UEs starting from the first subcarrier not used by *any* of the inner zones. Outer zone allocation occurs in ascending order according to data requirements, and is orthogonal between UEs. The second algorithm is identical to the first but takes frequency dependent fading into account in the outer allocation.

The paper claims that both algorithms improve the outage probability over a range of data rate targets. But there a number of flaws. First, the algorithms are biased toward inner zone UEs, since these UEs always get scheduling preference. This could lead to starvation of outer zone UEs. Second, the quantity of outer resources is determined by the largest inner zone requirement, which means that non-uniformity between cells will result in gross under-utilisation of resources.

In [113], an interference coordination mechanism is proposed whereby an overloaded cell communicates its status to its neighbours (specifically those neighbours adjacent to the regions of overload). The status includes the frequencies which are experiencing the most interference and the priority of the request. If an adjacent cell can spare the resource, either because it is lightly loaded or has low priority traffic that can be cut back, then the adjacent cell will reduce its transmit power on those sub-carriers indicated as having high interference by the communicated status report. This is illustrated in Figure 2-36 below. Note that following a subcarrier set power reduction, potentially all neighbours can increase their power over these subcarriers.

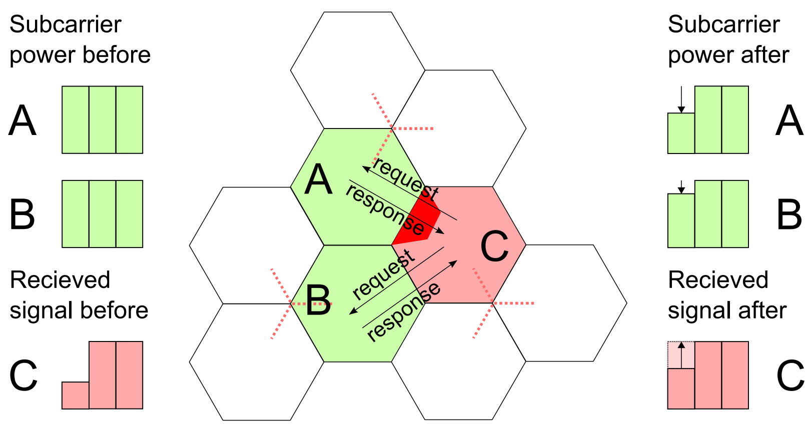

Figure 2-36: The diagram illustrates the transmit power on 3 sets of subcarriers for the cells A and B on the left and the poor signal on the first set of subcarriers for cell C. In the middle the overloaded cell C sends mute requests to cells A and B and they both respond. The effect of the mute requests are shown on the right where it can be seen that A and B have reduced their power on the requested subcarrier set, to the benefit of C obtaining a better received signal.

A simple simulation was performed to test the idea. With one overloaded sector being surrounded by "normally" loaded cells, the authors claim a 40% increase in sector throughput with marginal decrease in the throughput of neighbouring cells. The idea is an obvious one, and under the limited conditions tested appears to worthwhile. It is not however simulated with enough rigour to come to any solid general conclusions.

Distributed Solutions Without Communication

In [114] interference coordination is framed as a distributed *potential game* problem where the BSs represent decision makers whose goal is to locally minimise their experienced interference. A deterministic solution is derived which leverages a Hamiltonian tour of the cells to improve iteratively the potential function. A completely distributed probabilistic alternative follows in which all cells act simultaneously. Both strategies are based on the simple idea of greedy preferential selection of high SINR resources. Both schemes compare favourably to a globally derived optimal solution obtained using integer programming.

The results illustrate that a simple distributed algorithm can in principle obtain near-optimal inter-cell resource coordination in order to minimise interference. However, each cell is tasked to assign less than 1/3rd of the number of resource units, and user positions are not taken into account. This means that a prototypical reuse three solution can follow that has maximal utility. The efficacy of this solution is enhanced, because the work does not consider frequency selective scheduling. No information is provided as to how the algorithm will perform at higher loads, but since perfect coordination will no longer be possible, it seems almost certain that the performance of the approach will be reduced.

An exceptional study is presented in [115] for the downlink, in which two distributed algorithms are developed for computing transmit powers for dynamic soft reuse sub-bands. The first algorithm assumes communication among cells, and the second is a heuristic derivative that requires no inter-cell communication.

The concept is based on computing the gradients of sub-band transmit powers with respect to a global utility function. This information is used to make changes to transmit powers in the direction of gradient ascent. The gradients are estimated using information exchanged between neighbours: each cell tells its neighbours how much it causes them interference, for each sub-band. This is itself an estimate based on relative RSRP strengths and SINR measurements. Each cell then has an idea how it affects its neighbours, and uses this information to estimate the gradient of its transmit powers relative to the estimated global utility.

It is claimed that both schemes can improve the 5th percentile throughput for a given mean user or sector throughput, relative to reuse one. It is also claimed that the algorithm can track changes in UE distributions.

With respect to the work in Chapter 4 one might imagine that the algorithm is able to tune soft reuse to the ever-changing dynamics of the network, and mitigate the problems identified with static solutions. However, the results could seem somewhat optimistic. The work uses an unbounded Shannon formula to compute throughput. Shannon tends to be a little steeper than realistic mappings such as that shown in Figure 6-3, especially toward the higher and lower ends. This could lead to an overestimation of the effect of SINR improvements.

Another point worth examining is that the fast fading channel assumes a velocity of 20km/h, this is on the high side and might underestimate the potential of frequency selective scheduling. And indeed, in another work by the same group [116] they concede that frequency selective scheduling itself performs a kind of interference coordination and state that "the benefit of FFR in the case of fast fading is reduced or - in extreme cases - may be even non-existent". So there is some doubt about the efficacy of the approach in practice.

The same group has produced a similar scheme for the uplink and a general discussion of the work of the group can be found in [117]. This is titled to suggest it is LTE specific, but the downlink results presented are for the same generic OFDMA system used in [115].

Discussion

Opportunistic Scheduling It is surprising how often works compare their proposed solutions to random [94, 107, 99] or other non-realistic schedulers [96, 100]. This limits the extent to which conclusions can be drawn, especially in light of work which demonstrates the ability of the scheduler alone to manipulate cell-edge throughput [97, 98, 118, 119].

This is perhaps the most important take away message from this literature review: research that produces a novel interference mitigation strategy, if it is desired that it be considered seriously by mobile phone operators, must be compared with a realistic sector-independent scheduling strategy under reuse one, as this is the manner in which extant networks are actually being deployed. However, even outside of the industrial context, this practice should be encouraged, because as is observed in [116] "some degree of inter-sector interference avoidance is present "automatically"" with channel-aware scheduling. Thus, channel aware scheduling provides a much more sensible and challenging competitor than random scheduling for performance comparisons.

Despite this perspective, some conclusions *can* be inferred from the body of work. First of all, if frequency selective scheduling is unavailable, then enforcing orthogonality appears beneficial when the load is low enough to reduce resource collisions substantially. Randomised resource usage is otherwise the preferred strategy. In addition to being inferred here, this position is supported explicitly in [120] and [121].

Furthermore, an inverse relationship between load and technique success seems apparent and is supported in [109], where gains were seen at low loads but not at high loads, in [111] where the gains diminished with increasing load, in [110] where gains occurred over a larger range of SINRs at low loads than they did at high loads, in [114] where it was speculated that the gains seen would diminish at high loads, and by the work in Chapter 5 of this dissertation. So more generally, it seems that it is easier to obtain gains from interference coordination at low loads due to the greater flexibility available. This is not particularly surprising, but it is useful to observe that simulation work supports the intuition.

Inter-cell Fairness A common question which arises from [99, 100, 101, 103, 104] and to a lesser extent [94, 97, 96] is inter-cell fairness. It is well known that for localised scheduling decisions, a maximum C/I scheduler (See Section 2.1.6) is not representative of practical implementations because it favours particular UEs at the expense of others. With the exception of [95, 109], none of the above cited work provides comment on the fairness of their promoted schemes with regard to inter-cellular differences in throughput. And since it is in theory possible that large mean total and 5th percentile throughputs could appear as an emergent property of great variance in cell throughputs, care should be taken to assess this.

The algorithm in Chapter 5 of this dissertation supports fairness by using the same evaluation metric in each of the communicating cells. And since the metric considers the UEs mean rate and instantaneous conditions, large inter-cellular differences should be reduced as a consequence.

Resource restrictions The general approach taken by most dynamic interference mitigation schemes, is to place restrictions on the use of resources in adjacent cells. The motivation for this can only be that the benefits of these restrictions outweigh the shortcomings. In some cases this is taken into consideration explicitly [99], in others the benefits of doing so are left unjustified albeit demonstrated true against a random scheduler [95, 96], and finally in some cases, heuristics are used to decide whether or not to enforce orthogonality between UEs, without *any* justification [101, 103, 104, 122]. The use of unjustified heuristics is a weak point in the literature, because it does not provide any understanding of *why* a given technique works or *how* its success can be generalised.

Chapter 4 in this dissertation opens by quantifying the SINR gains that might be expected from the approach presented therein, and Chapter 5 in this dissertation extends this analysis by examining the geographic likelihood of obtaining a benefit from enforced intrasite resource orthogonality. Admittedly, it is easy to overlook the need for such analysis, when a promising technique nevertheless shows gains over its competitors. But it would be prudent if future research in this domain placed a greater emphasis on analysis and explanation.

A final point worth mentioning about resource restrictions is that, as observed in Chapter 5 of this dissertation, the gains from frequency selective opportunistic scheduling are large enough that placing restrictions on resource allocation, and thus on the flexibility of opportunistic scheduling, is a strategy that pays off only in particular circumstances. This is why, as seen in Chapter 4 of this dissertation, purely static schemes offer limited value and are unsuitable for heterogeneous deployments. A similar conclusion is reached in [117] where the following is said of static schemes "For situations where FSS (Frequency Selective Scheduling) gains can be exploited, there is no ICIC gain over FSS with universal reuse".

Motivation for a novel idea In this review of dynamic schemes, many techniques have been presented that promise excellent gains. The most promising of these is probably [115]. But like the others, the main problem with this approach is that the vendors will never actually implement it on their hardware, because doing so would remove their competitive advantage, which traditionally has come from (often patented) proprietary differentiation.

Note that this is not to discount the good ideas they produce. Two worth mentioning are: (i) the idea of focusing only on one or two dominant interferers or restricting the spatial extent of coordination [99, 101], backed up by empirical results which suggest that almost all the gain can be obtained by coordinating with only a few cells, (ii) the use of RSRP to estimate the component contributions of the interference element of SINR [115, 101].

For these reasons, in Chapter 5 of this dissertation it was decided to focus on intrasite scheduling because: (i) it is easy to do. (ii) it is easy to analyse, (iii) it is a relatively inexpensive, and should it offer reasonable gains, the vendors may be convinced, not to adopt the exact algorithms proposed, but at least to pursue the idea itself.

Probably the closest work to that of Chapter 5 of this dissertation is found in [123]. Interestingly the authors point out the deterioration in throughput that occurs because of the mismatch in measured and actual CQI, due to on-off transmissions. Four schemes are compared, the first just obtains intrasite orthogonality, the second randomly selects a set of resources and sticks to them, the third uses frequency selective scheduling but with a correlation so that picked resources are picked again with a greater probability, and the last picks resources based on CQI averaged over all UEs. Note that there is no comparison with purely opportunistic scheduling but the results show that the quality metric-based approach performs best, which is consistent with the results in Chapter 5 of this dissertation.

2.4 Chapter Summary

In this chapter three areas of the literature have been reviewed. In Section 2.1 general background information related to mobile phone networks and technology was presented, to ground the subsequent discussions for non-experts.

Section 2.2 presented an overview of Mobile WiMAX technology in order to provide system understanding for the work in Chapter 3, which exploits the technology. The concept of frequency selective scheduling was introduced, and the work to date in examining the impact of velocity on its performance reviewed. In summary, the motivation for the work in Chapter 3 was attributed to two shortcomings of the existing work: the lack of information across the whole velocity range, and the existence of conflicting results.

In Section 2.3 a comprehensive review on interference avoidance techniques was presented. The subsection on static schemes described how the literature contains many conflicting results, which are sometimes presented without evidence, or sufficient justification. This motivated the need for the more rigorous analysis, and elucidation of determinant factors, presented in Chapter 4. The subsection reviewing dynamic schemes, illustrated the areas where further work was deemed to be required. The business constraint for simple solutions suitable to extant network equipment was identified, and motivated the case for the interference mitigation approach presented in Chapter 5.

The next three chapters contain the research contributions of this dissertation: portfolio reports which address the problems identified in literature just reviewed. Chapter 3 examines the impact of MS velocity on the performance of frequency selective scheduling, Chapter 4 analyses static soft reuse for interference mitigation and derives the determinant performance factors, and Chapter 5 develops and tests a dynamic intrasite scheduling algorithm that addresses the problem of inter-cell interference.

Chapter 3

The Impact of MS Velocity on the Performance of Frequency Selective Scheduling in IEEE 802.16e Mobile WiMAX

technical report

vodafone

The impact of UE velocity on the performance of frequency selective scheduling in IEEE 802.16e Mobile WiMAX

Reference:	**Filed as**
VFGR&D_2011_1234	Sharepoint
Version	**Status**
1.0	FINAL
Date	**Approved**
04/01/2011	David Lister
Author	**Date/Signature**
Ashley Mills, David Lister, Prof. Yusheng Ji, Dr. Marina De Vos	**4-Jan-11** …………………………………………………………………………………………
Department	**C2 – company confidential**
Group R&D .UK	

Scope

This document describes how the benefits of frequency selective scheduling are dependent on UE velocity. Estimated upper bounds on performance for different velocities are given. Recommendations for practical deployments are provided based on these estimates.

Document history	Date	Reason
1.0	30-09-08	First release

technical report

vodafone

Document distribution

Mike Walker, Group R&D Director		
Trevor Gill David Lister Luke Ibbetson Ralf Irmer Adam Pollard Alison McIntyre	Kenny Graham Andrew de la Torre Steve Allen Peter Cosimini	Group Technology
Tim Frost, Group R&D Standards Prakash Bhat Soon Leh Ling		

Table of Contents

Executive Summary

Mobile WiMAX 802.16e and LTE use the OFDMA air interface. The OFDMA air interface allows resources to be allocated in blocks that each span a fraction of the total number of subcarriers and temporal units belonging to a frame. The subcarrier groups which comprise each of these blocks are received with different strengths by different users, depending on the location of these users relative to the base-station and their mobility. Thus each user experiences a personalized fading profile over all available resource blocks and receives some resource blocks with greater strength than others. The scheduler, to the best of its ability, can take advantage of this fact and shuffle the resource allocation such that each user is allocated the resource blocks which they receive with greatest signal strength. This is called frequency selective scheduling.

Frequency selective scheduling relies on UE reports concerning the channel quality of available subcarriers and as such its success is dependent on the temporal stability of the channel at the subcarrier level and the timeliness and accuracy of channel quality reports.

When a UE is moving, the UE experiences a channel with reduced temporal stability, in comparison to a stationary channel. Thus it is commonly accepted that the efficacy of frequency selective scheduling will be reduced in this situation.

In the extreme, when CQI reports are completely inaccurate and outdated, it is senseless to use frequency selective scheduling since choosing a resource block using this information is the same as choosing a random resource block. Since the subcarriers in the resource blocks are adjacent and will fade dependently, choosing resource blocks randomly will cause an increase in block errors. In this situation it is therefore more sensible to spread the resource block over a frequency diverse set of subcarriers in order to benefit from frequency diversity and reduced block error.

This work examines the area in between the two extremes; a stationary and predictable channel in which frequency selective scheduling can be used to advantage; and a highly dynamic and unpredictable channel in which frequency diverse scheduling can be used to advantage. We ask at what UE velocity or detriment to channel conditions is it optimal to switch from frequency selective scheduling to frequency diverse scheduling.

This is achieved by examining the performance of frequency selective scheduling and frequency diverse scheduling, in a simulated Mobile WiMAX environment, for different UE velocities. Both link and system level experiments are executed. It is found that at low velocity, frequency selective scheduling outperforms frequency diverse scheduling. This performance steadily declines with increasing velocity up until between 15 and 20kmh, frequency diverse scheduling then outperforms frequency selective scheduling.

The conlcusion is that frequency selective scheduling is a viable means of increasing throughput in WiMAX and throughput can almost be doubled at low velocity under ideal conditions. A frequency selective scheduling scheme such as the WiMAX Band AMC mode of operation should be selected for urban cells in which the traffic source is dominated by pedestrians or stationary users. In rural environments, in which a system is deployed specifically to serve roads or railways, a frequency interleaving approach such as the PUSC mode of operation should be selected.

1 References

[1] Krishna Balachandran and Doru Calin and Fang-Chen Cheng and Niranjan Joshi and Joseph H. Kang and Achilles Kogiantis and Kurt Rausch and Askok Rudrapatna and James P. Seymour and Jonqyin Sun. Design and analysis of an IEEE 802.16e-based OFDMA communication system. *Bell Labs Technical Journal*, 11:53-73, 2007.

[2] John G Proakis. *Digital Communications*. McGraw-Hill Science/Engineering/Math, Forth edition, 2000.

[3] 3GPP TS 36.300 - V8.5.0 - Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Overall description; Stage 2 (Release 8). 2008.

[4] Intel Corporation. 3GPP TSRG1-01-0030, Further Results on CPICH Interference Cancellation as A Means for Increasing DL Capacity.

[5] Riato, N. and Sorrentino, S. and Franco, D. and Masseroni, C. and Rastelli, M. and Trivisonno, R. Impact of Mobility on Physical and MAC Layer Algorithms Performance in Wimax System. *Personal, Indoor and Mobile Radio Communications, 2007. PIMRC 2007. IEEE 18th International Symposium on*, :1-6, 2007.

[6] P802.16Rev2/D4 April 2008, DRAFT Standard for Local and metropolitan area networks Part 16: Air Interface for Broadband Wireless Access Systems. Consolidation of IEEE 802.16-2004, 802.16e-2005, 802.16f-2005, and 802.16g-2007..

[7] WiMAX Forum. WiMAX Forum Mobile System Profile Release 1.0 Approved Specification (Revision 1.4.0: 2007-05-02).

2 Abbreviations

3GPP	Third Generation Partnership Project
AWGN	Additive White Gaussian Noise
BAMC	Band AMC
BS	Base-Station
CDF	Cumulative Distribution Function
CQI	Channel Quality Index
HARQ	Hybrid ARQ
ITU	International Telecommunications Union
LTE	Long Term Evolution
MCS	Modulation and Coding Scheme
MIMO	Multiple Input Multiple Output
OFDMA	Orthogonal Frequency Division Multiple Access
PER	Packet Error Rate
PUSC	Partial Use of Subcarriers
QAM	Quadrature Amplitude Modulation
QPSK	Quadrature Phase Shift Keying
SINR	Signal to Interference plus Noise Ratio
SNR	Signal to Noise Ratio
UE	User Equipment
WiMAX	Worldwide Interoperability for Microwave Access

3 Introduction

OFDMA air interface allows data to be encoded over resource blocks that span arbitrary regions of space, time, and frequency. Figure-1 illustrates this. The Figure shows four spatial planes $(s = \{A, B, C, and\ D\})$ corresponding to four different propagation paths or antennas, five successive temporal units $a_{*,1}...a_{*,5}$, and 3 frequency subcarriers $a_{1,*},....a_{3,*}$. The lowercase a means that these elements belong to its uppercase counterpart, spatial plane A. Each element $s_{i,j}$ encodes one symbol, for example, a QPSK symbol.

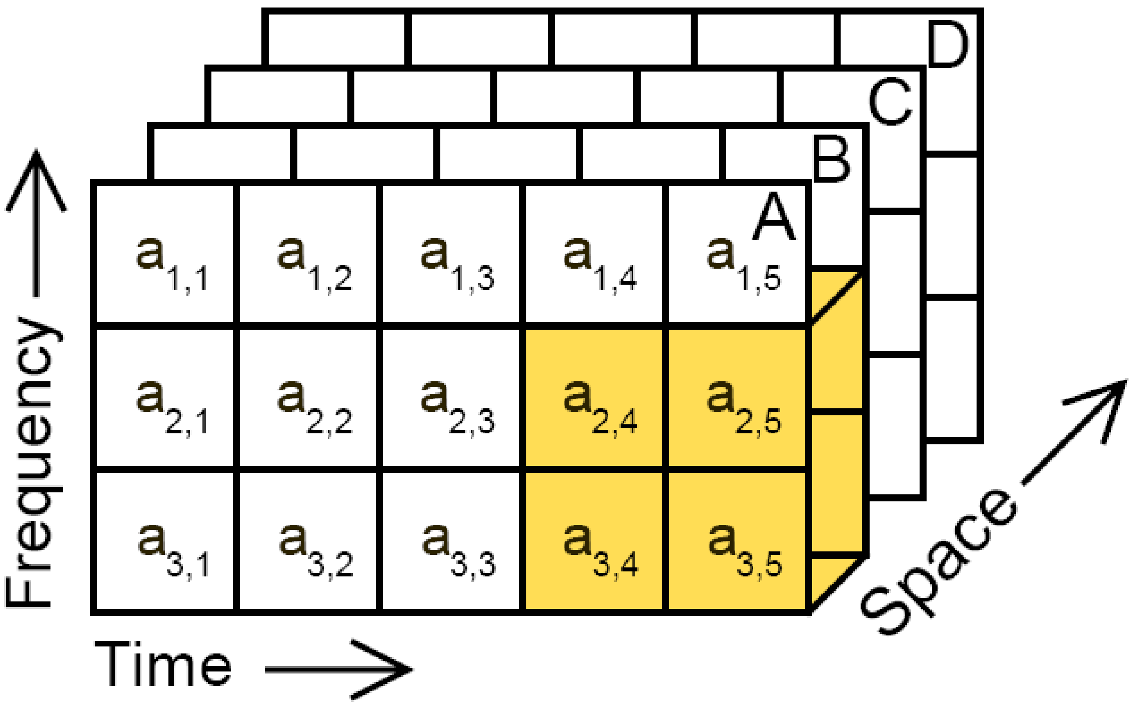

Figure 1: See text for description

The shaded area illustrates the concept of an arbitrary allocation of symbols to construct a data resource block. The resource block contains resource elements $a_{2,4}, a_{2,5}, a_{3,4}, a_{3,5}, b_{2,4}, b_{2,5}, b_{3,4}, and\ b_{3,5}$. The resource block spans two spatial channels A and B, two subcarriers, and two temporal units.

To simplify matters, in this work only a single spatial plane is considered. What we are concerned with is resource allocation; that is, how to construct resource blocks and how to allocate them to users. We are interested in how different ways of doing this might be suited to different radio scenarios associated with different users, specifically, we are interested in changing the way that resource blocks are constructed, based on the velocity of the user element.

There exist two fundamentally opposed approaches to doing this. Firstly, in what is called frequency interleaving, resource blocks are constructed by taking subcarriers drawn pseudo-randomly in such a way that they are spread approximately evenly over the frequency domain. In the second approach, resource blocks are constructed by taking subcarriers that are physically adjacent. We will compare these two approaches to resource block construction for users with different velocities as we believe that the former way is more suited to high velocity users and that the latter can be used to get larger gains from slow users. The motivation for this hypothesis will be explained next.

For a given OFDMA (Orthogonal Frequency Division Multiple Access) radio channel, let p be the probability that a subcarrier selected at random is subjected to interference above some threshold required for proper detection. Given N subcarriers, with successive frequency spacing greater than the coherence bandwidth of the channel, the probability that they will all simultaneously be subjected to such interference is p^N, because the channels fade independently. [2]

Whilst this evidently is not a comprehensive proof, it demonstrates intuitively why a signal, replicated across multiple subcarriers, is more likely on average to be received with at least one copy intact if the subchannels are spread in frequency in this manner than if the replicated signal is sent over adjacent subchannels that fall within the channel coherence bandwidth. The former concept is called introducing frequency diversity, or frequency interleaving, and is an often used approach to mitigate the localized nature of frequency selective fading in the absence of channel information.

Suppose however that, for a given UE (User Equipment), the channel quality of successive groups of adjacent subchannels can be estimated reliably. Then, assuming stability of this channel over the period of data transmission, relative gains can be obtained when sending data to this UE by choosing the subchannels which it receives with greatest signal strength. This is called frequency selective scheduling.

As the velocity of a receiver or transmitter is increased, the radio channel becomes more variable. This is illustrated in Figure-2.

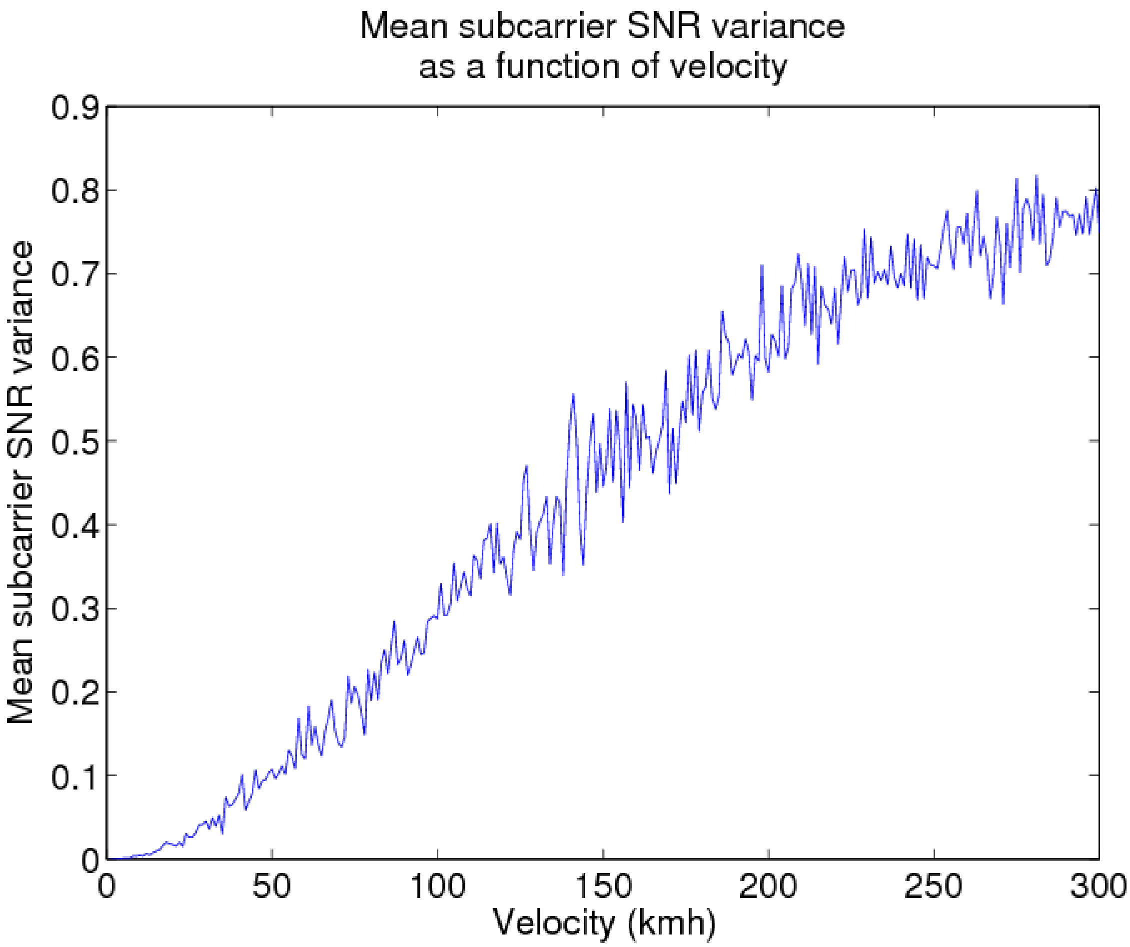

Figure 2: A Rayleigh fading channel with a Vehicular A tap delay and gain profile was used to generate fading envelopes for 108 physically adjacent subcarriers for a duration of 1.8ms. Subcarrier powers were sampled every 100 microseconds to obtain a total 18 samples. This resulted in an 108x18 matrix. The mean row variance (mean subcarrier variance) was computed and recorded. The velocity (via the Doppler factor) of the channel was varied from 0 to 300km/h in 1km/h steps and the sampling process repeated for each step. Averaged over 30 runs, the figure shows the mean subcarrier variance as a function of velocity.

Figure-2 shows the variance of subcarrier noise in a Vehicular A channel for different Doppler factors (velocities). The variance here provides a crude measure of the coherence of the simulated burst.

As the channel becomes more variable, intuitively it seems likely that the benefits of frequency selective scheduling will diminish, as the channel estimates taken at one instance will become less correlated in time with the next.

The hypothesis examined here is this ``common wisdom'' that frequency selective scheduling will outperform frequency interleaving at low velocities and only at low velocities. It is postulated that as the UE velocity is increased, the relative performance gain of frequency selective scheduling over interleaving will diminish and at some point become inferior on account of the diversity gains that make frequency interleaving robust to channel variations. Empirical validation and the optimal switch-point associated with this hypothesis are issues which have not been treated with significant rigor.

Here the performance of the two approaches; frequency selective scheduling and frequency interleaving are compared at different velocities using, next generation OFDMA technology, Mobile WiMAX, as a platform for investigation.

Mobile WiMAX has two subcarrier permutation schemes which are analogous to frequency selective scheduling and frequency interleaving; these are called Band AMC (Band Adaptive Modulation and Coding) and PUSC (Partial Use of SubCarriers), respectively.

Since OFDMA in Mobile WiMAX is a relatively new application, there are only two precedents to the question addressed here.

In [1], the system capacity and over the air performance of Mobile WiMAX was compared for various parameter configurations, including the use of PUSC and Band AMC permutation schemes. Two velocities were compared for each of these schemes; Pedestrian B at 3km/h, and Vehicular A at 30km/h. It was found that Band AMC performance exceeded PUSC cell throughput at both 3km/h (by 18%) and 30km/h (by 5%). Since no other velocity classes were tested, it is not clear at what point, if at all, the trend reverses.

This result is contradicted in [5] where a 50% performance gain for Band AMC over PUSC claimed for 3km/h but at the next velocity point tested, 8km/h, there is no gain and PUSC is claimed to outperform Band AMC by 5%. From the graphs provided, a crossover point of 7km/h can be estimated.

In the latter study, permutation specific schedulers are designed and tested for each of the two permutation types; PUSC and Band AMC. The schedulers both differentiate between QoS classes and employ conservatism in MCS selection and resource allocation when dealing with data with real time requirements. In addition, the rate of channel quality feedback is adjusted dynamically per user according to their estimated velocity class; slower UEs report less frequently. And it is demonstrated that uplink savings can be made by doing so, although at the expense of channel estimation accuracy and consequently downlink throughput. Whilst interesting, use of these complicated techniques and the multi-faceted report that consequently emerges, clouds the issue of interest here; namely, how the performance of frequency selective scheduling varies as a function of velocity.

Therefore, given the discrepancy between these two results, and the complexity which clouds the latter, it is pertinent to repeat the experiment in a more focused manner with regard to system setup and over a larger range of velocities.

To re-iterate, In this work therefore we compare the performance of frequency interleaving with frequency selective scheduling at different velocities to determine the practical crossover point for normal operation.

In Section-2 link level simulations are performed and the results analyzed, in Section-3 the system level simulations are focused on. A brief discussion is given in Section-4 and the paper concludes in Section-5.

4 Link level simulations

4.1 Overview

The purpose of these experiments was to compare the performance (throughput for a given SNR) of frequency interleaving (PUSC) with the performance of frequency selective scheduling (Band AMC) for different velocities, in order to determine a practical crossover point.

This is a two stage process; it is computationally impractical to fully simulate frame transmission at the system level because thousands of frame transmissions are required per second, and each transmission is computationally expensive. Link level simulation is an abstraction that reduces a frame transmission computation to a lookup operation. This is achieved by performing offline computation of the SINR to PER relationship for a large number of characteristic situations and SINR values. These lookup tables can then be used in a system level simulation. The link level results themselves also have meaning.

The first step is therefore to perform link level simulations.

4.2 Procedure

Link layer simulations were executed to test the performance of Band AMC and PUSC at different velocities. The parameters shown in Table-1 were used.

Index	Parameter	Value
1	Link direction	Downlink only
2	Carrier frequency	2.6 GHz
3	Bandwidth	10 MHz
4	NFFT size	1024
5	Frame length	5ms
6	Cyclic prefix length	1/8th
7	Forward Error Correction (FEC) type	CCTB
8	Slot type when Band AMC is used	2x3
9	Pilot boosting	2.5db
10	CCTB circulation depth	60
11	CCTB decision depth	60
12	Block height	6 Subchannels
13	Block width (Band AMC)	18 OFDMA Symbols
14	Block width (PUSC)	12 OFDMA Symbols
15	CQI feedback period	1 frame
16	Channel estimation assumption	Perfect
17	Channel type	Jakes, VehA

18	Number of transmit antennas	1
19	Number of receive antennas	1
20	MIMO	None
21	HARQ	Disabled

Table 1: Link level parameters

Parameter 1 indicates that the simulation is downlink only. Parameters 2 to 8 are taken from The WiMAX forum Mobile system profile, Release 1 [7], this is a restriction on the 802.16e parameter space that provides a standard to which manufactured devices must conform if they wish to be WiMAX Release 1 certified. Parameter 9 is as specified in the 802.16e standard [6].

Parameters 12, 13, and 14 define the size of the transmission burst to simulate. The rectangular frequency-time shapes for transmission bursts for Band AMC and PUSC in this implementation are different: a Band AMC block is 50% longer in time than a PUSC block (18 vs 12 OFDMA symbols), whereas a PUSC block takes up 50% more frequency spectrum than an Band AMC block (144 vs 108 subcarriers). This is because Band AMC and PUSC subchannels have lengths of respectively 3 and 2 OFDMA symbols and widths of respectively 18 and 24 subcarriers. Further, in order to use a repetition rate of 6 using our implementation the number of subchannels had to be a multiple of 6. For comparable throughput results however, the total number of resource elements in each type of burst needs to be equal. The settings here ensure this, giving a burst size of 36 slots.

The shape and size of a Band AMC burst is important because the decoder performance depends on how coherent the burst is with respect to deviations in SINR. In the absence of quantitative simulation results it is difficult to speculate on the optimality of the choice here with regard to each encoding type Band AMC and PUSC. And in fact such an optimality study could constitute an independent thread of research. Therefore the results should be considered with this in mind.

Parameters 15 and 16 specify optimistic feedback requirements and channel estimation assumptions. These were chosen to provide an upper estimate of performance.

Parameter 17 indicates that the channel used was Jakes ITU Vehicular A. This is a simple multi-tap channel where the delays and mean powers are taken from a known reference [4] to facilitate comparison with other results.

MIMO is disabled and no receive diversity is used, as indicated in 18,19,20, and 21.

Simulations were executed for the frame encoding types shown in Table-2:

Permutation scheme	Description
PUSC	PUSC encoded subchannels
BAMC 1	Band AMC with subchannels 1-6 always selected from an SINR ordered list sorted in descending order. (Best set)
BAMC 2	Band AMC with subchannels 7-12 always selected from an SINR ordered list sorted in descending order. (Second best set)
BAMC 3	Band AMC with subchannels 13-18 always selected from an SINR ordered list sorted in descending order. (Third best set)

Table 2: Description of link level permutation schemes tested

The best set, second best set, and third best set Band AMC variations are intended to emulate the action of frequency selective scheduling to differing degrees of benefit. The PUSC variation is used to emulate the action of spreading.

The following velocities were studied, this range covers with finer granularity the the ranges covered by both [5] and [1] and considerably extends the range of the latter in order to see how far the benefits of frequency selective scheduling extend.

$$\{0.1,1,2,3,4,5,7,10,13,15,20,25,30,40,50,60,70,80,90,100,110,120,140,160,180,200,250,300\}$$

The following MCS schemes were examined:

{ QPSK $\frac{1}{2}$ Repetition rate 6, QPSK $\frac{1}{2}$ Repetition rate 3, QPSK $\frac{1}{2}$, QPSK $\frac{3}{4}$, 16 QAM $\frac{1}{2}$, 16 QAM $\frac{3}{4}$, 64 QAM $\frac{2}{3}$, 64 QAM $\frac{3}{4}$ }

These MSC schemes were chosen the same as used in [1] in order to facilitate better comparison of results.

For each frame encoding, MCS, and velocity, SNR values between -20 and 30 in steps of 2 dB were examined. For each SNR point, 300 blocks were simulated and the average PER recorded.

4.3 Results

This gave a set of SNR vs PER curves for each combination of MCS scheme, permutation scheme, and velocity. For illustration, an example set of curves is shown in Figure-3.

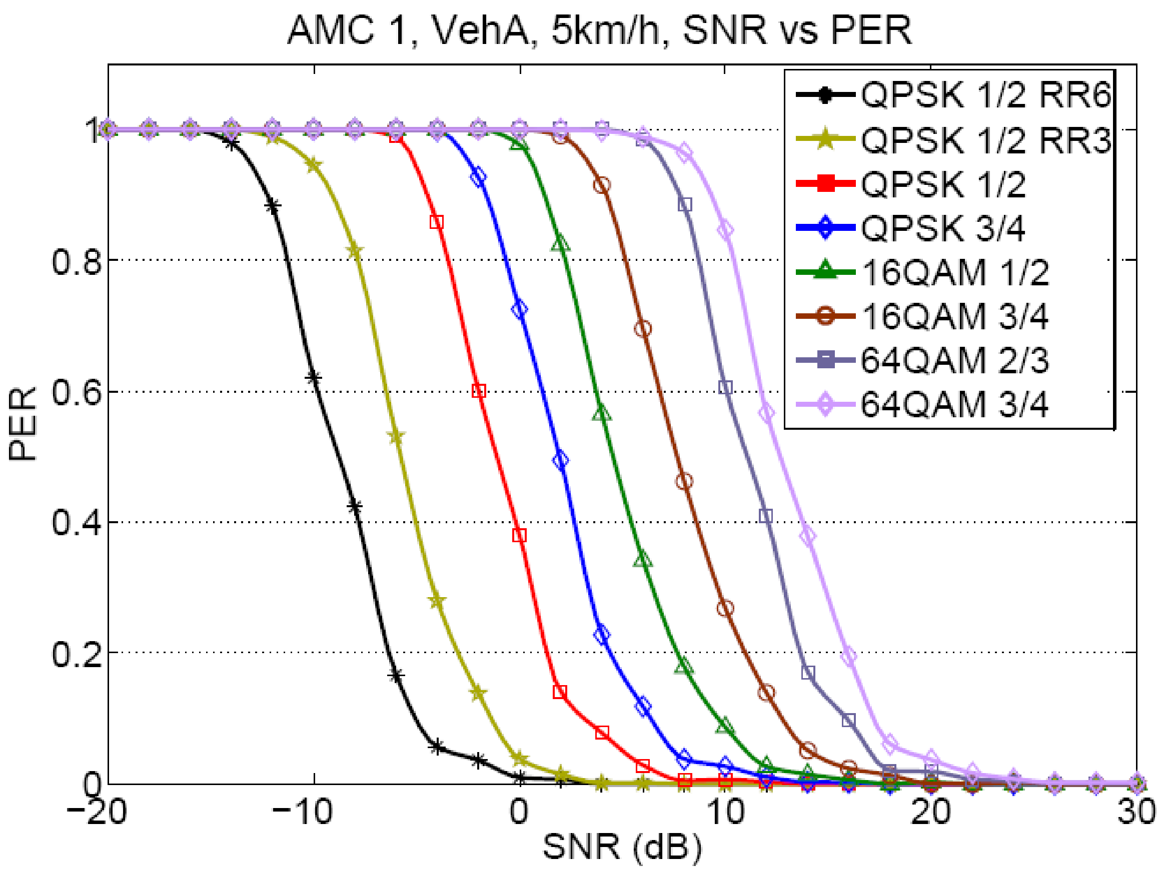

Figure 3: SNR vs PER for different MCS schemes for BAMC 1 encoding at 5km/h.

Given such a set of MCS curves, it is normal to derive a new function which provides, for each SNR point, the largest achievable throughput.

At a given SNR point, each MCS will have its own associated PER. The product of the theoretical throughput of that MCS and (1-PER) gives its achievable throughput. The maximum of these is the maximum achievable throughput at the given SNR.
While theoretically correct, practically such a function would be incorrect. The reason is that the system level will never select an MCS having PER greater than 10%, whereas the function described above, for a given SNR, could have its maximum achievable throughput using a higher order MCS that nevertheless has PER greater than 10%.

Thus, a new function is defined, which returns for each SNR point, the maximum achievable throughput that has PER less than or equal to 10%. This function is called a waterfall curve and is illustrated for VehA 5km/h with the AMC_1 frame in Figure-4.

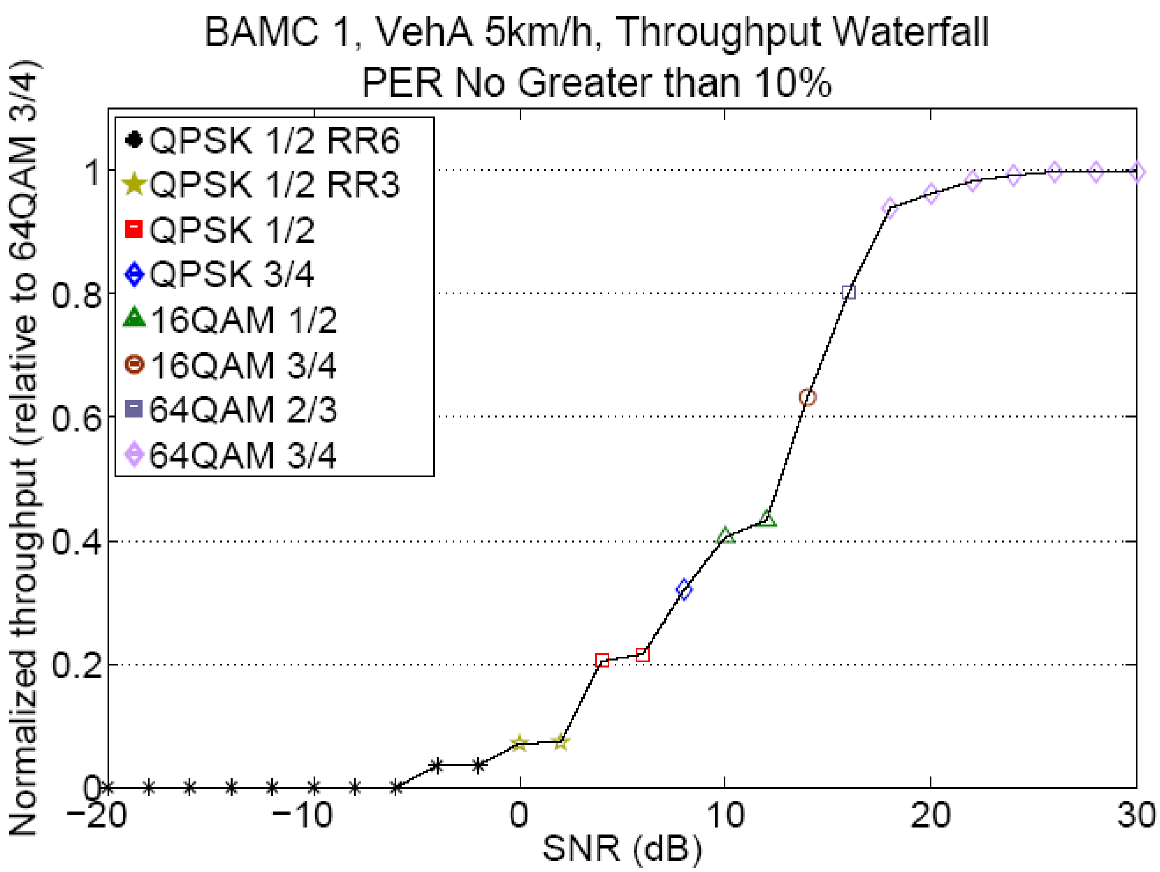

Figure 4: ``Waterfall curve'' for BAMC 1 encoding at 5km/h.

Given such a waterfall curve, it is possible to compute the mean throughput over all SNRs, to give a value which characterizes the link-level performance associated with the velocity and permutation scheme that determine the curve.

This procedure was performed for each velocity and permutation scheme, to give the results shown in Figure-5

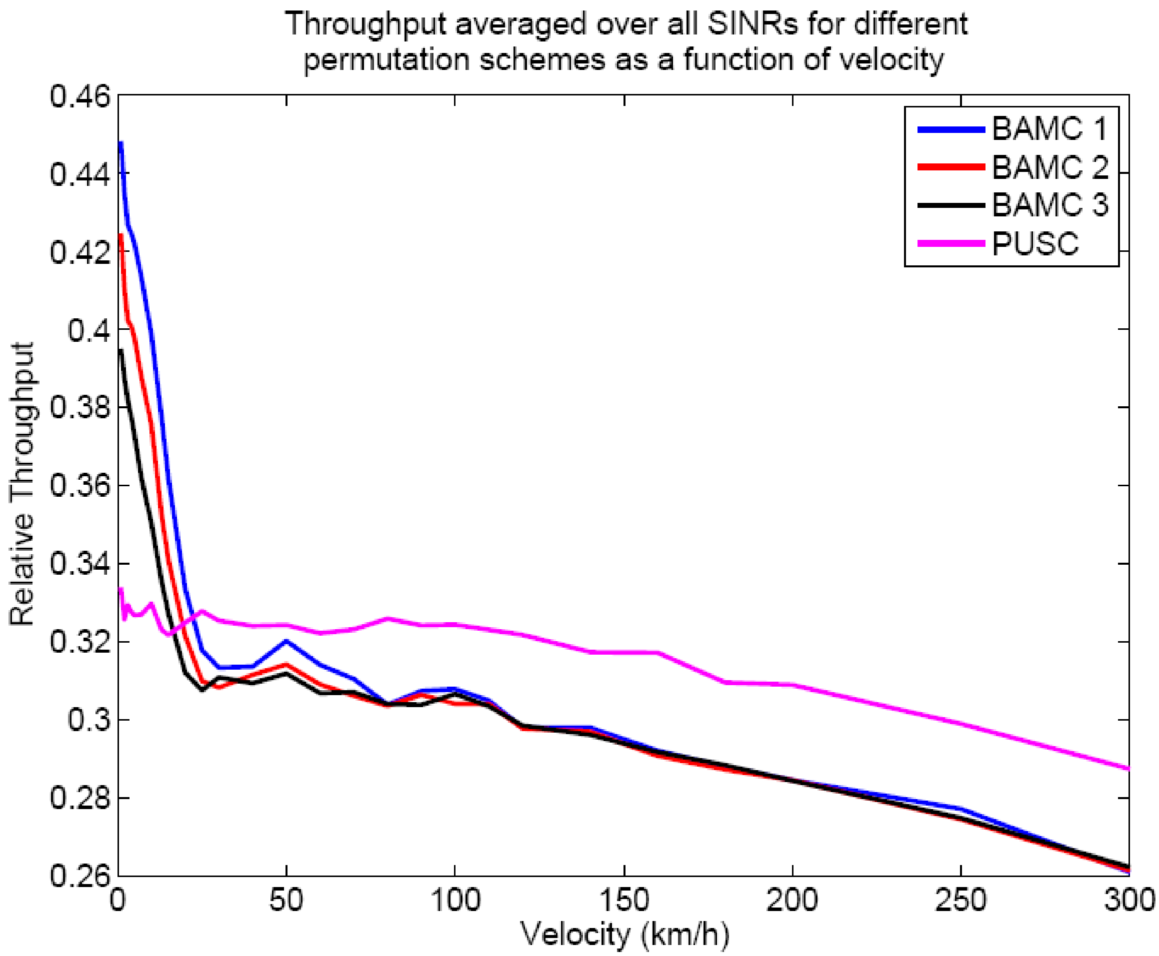

Figure 5: Link level performance results for the four different permutation schemes tested as a function of velocity.

These results confirm what in theory is predicted. At low velocity, it is possible to take advantage of frequency selective scheduling, because the temporal coherence is high. Also as expected, over this region, selecting the better subcarriers results in better performance. Figure6 the zooms in on the the first 30km/h for clarity.

At between 17 and 22 km/h, it no longer becomes efficient to use Band AMC because it results in insufficient spreading of the signal in a temporally uncorrelated environment. Groups of subcarriers suffer simultaneous deep fades to the detriment of performance.

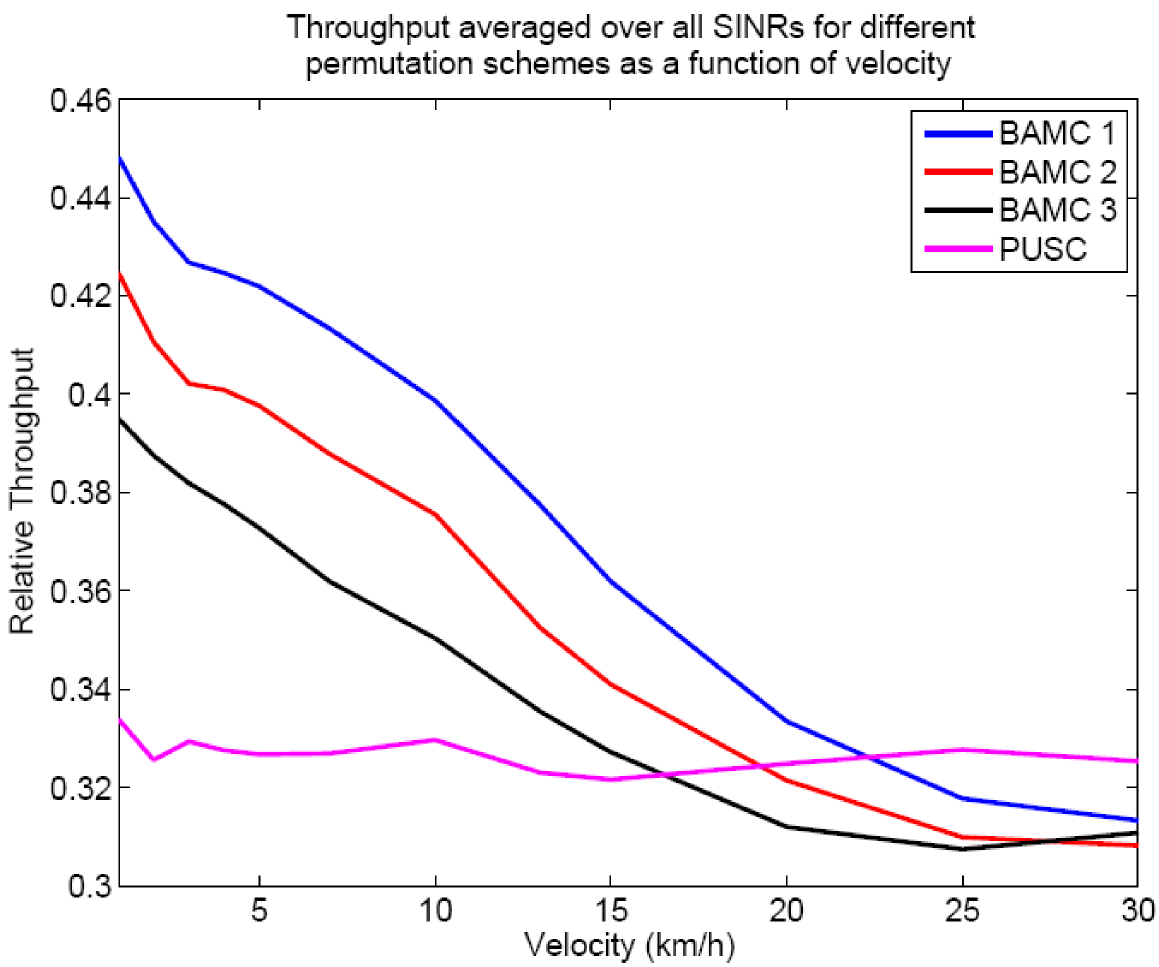

Figure 6: Link level performance results for the four different permutation schemes tested as a function of velocity. This plot only shows velocity upto 30km/h.

5 System level simulations

It is necessary to ask how these link-level results translate to performance gains at the system level.

5.1 Procedure

A full list of parameters is given in Table-3. The system layout was 1 tier containing 21 cells with wrap-around. Mobility of users was set according to the velocity examined. The range of velocities studied was the same for the link level case above. The scheduler used was simple Round Robin. The reason that this is acceptable, given that a frequency selective mode of operation was under test, is because the link level curves here already capture the effects of on average choosing the n th set of best subcarriers.

Parameter	Value
Simulation duration	10000 frames
Link direction	Downlink only
CQI averaging window	4 frames
Fading granularity	10m
AMC threshold	10%
Site to site distance	1km
Sector count per site	3
UE per sector	25
HARQ	Disabled
Handover	Make before break
Traffic model	Full queue
Scheduler	RoundRobin
Reuse	1

Table 3: System level simulation parameters

The site to site spacing ensures that coverage exceeds thermal noise and that the system is interference limited.

A comment is necessary here with regard to the system level results which follow.

The system level simulator does not support Band AMC style CQI measurement and as such it is not possible to tell from the simulation whether Band AMC CQI measurements would have higher values than the preamble (i.e the PUSC) CQI measurements. The results therefore reflect only a difference in the lookup curves, i.e the link level curves and hence are only suggestive of possible throughput values and gains.

5.2 Results

The results are shown in Figure-7:

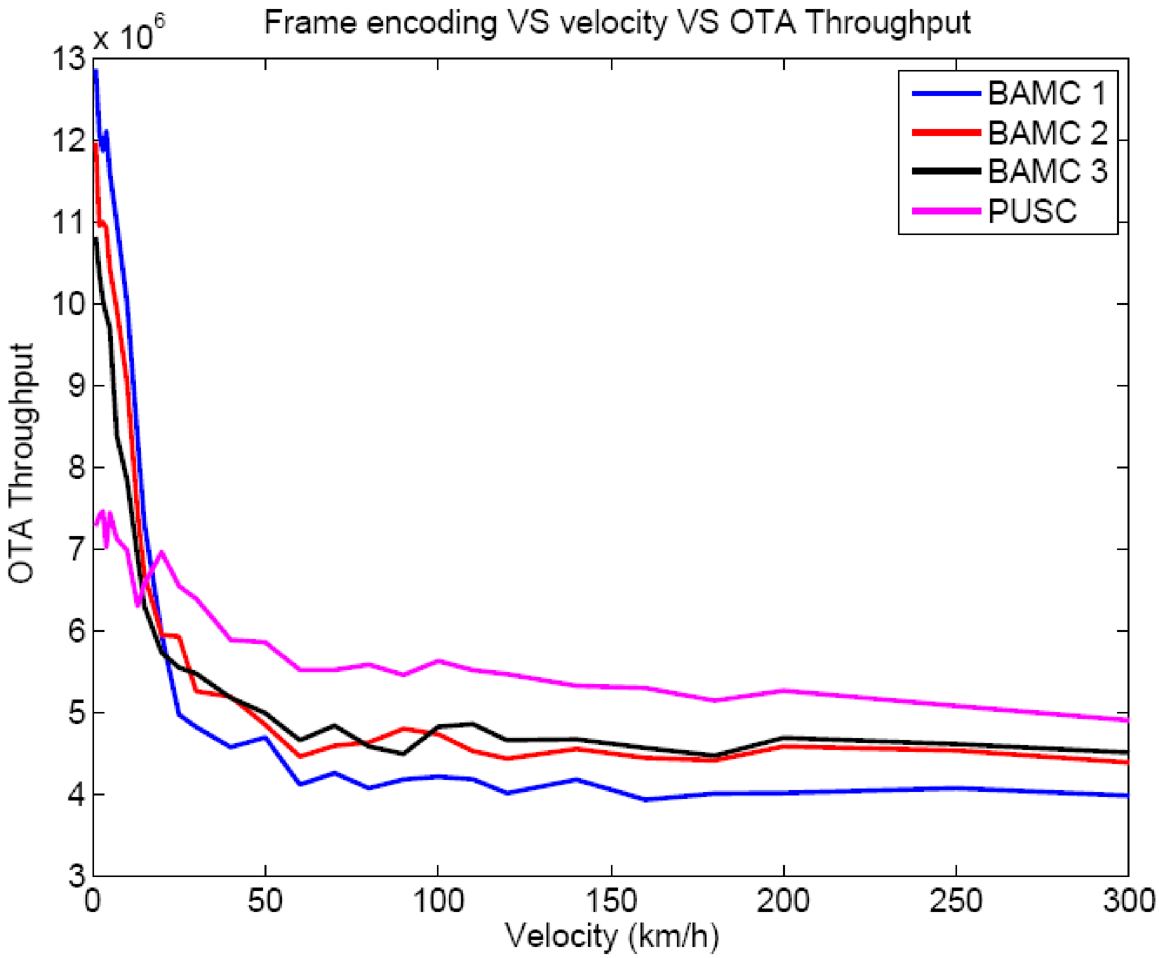

Figure 7: System level performance results for the four permutation schemes tested, as a function of velocity.

Figure-8 shows the results from Figure-7 but magnified to show only the first 30km/h.

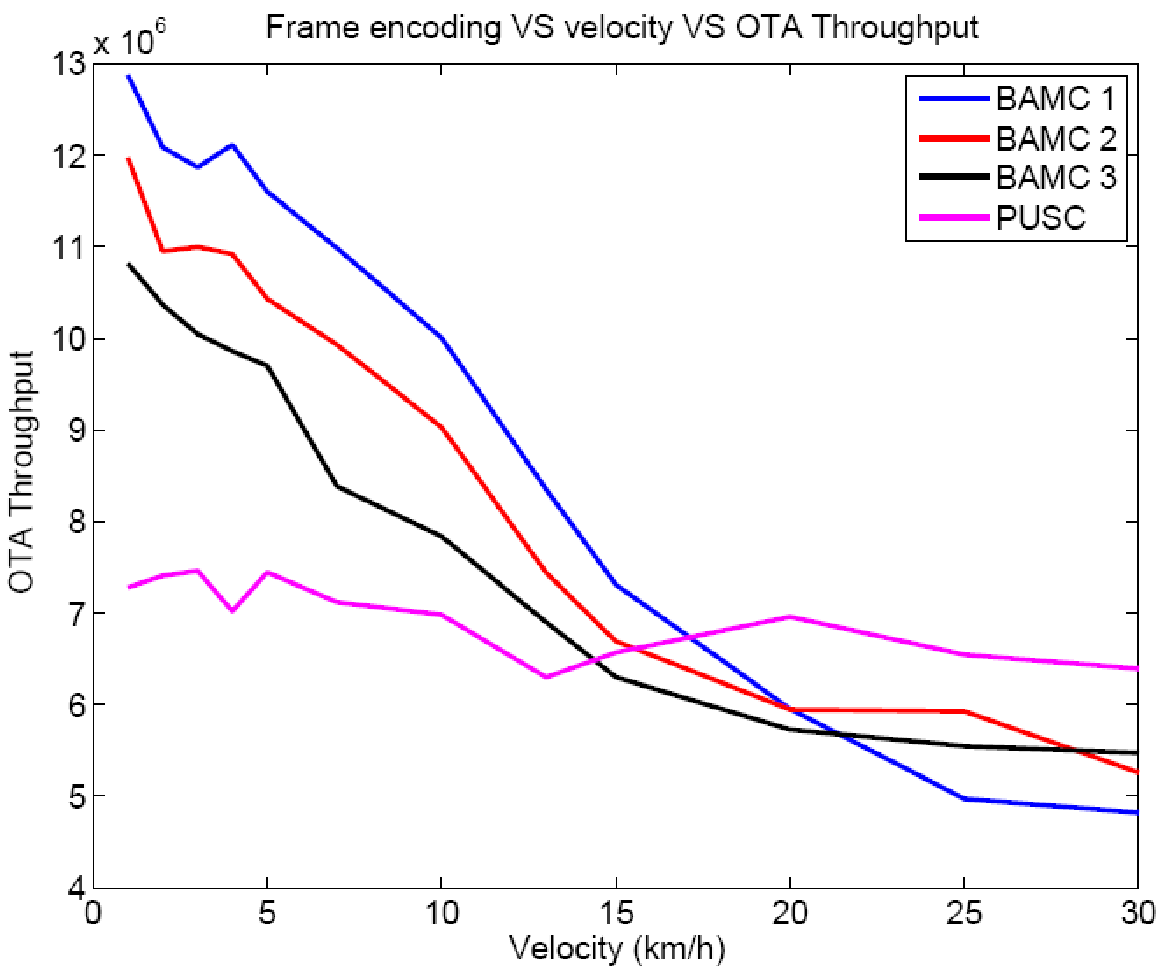

Figure 8: System level performance results for the four permutation schemes tested, as a function of velocity. This plot only shows velocity upto 30km/h.

At 1km/h the difference between AMC_1 and PUSC is very large, almost twice the over the air throughput (7Mbps vs 13Mbps).

The crossover point appears to be at approximately 15km/h. With reasonable confidence it can be claimed that there is still a benefit to be had at 10km/h.

6 Discussion

Frequency selective scheduling is a method to take advantage of known channel information whilst it is still valid. The validity of this information was expected to degrade at higher velocities and this has been demonstrated, evidenced by reduced performance.

The following things could be the subject of future inquiry:

• Burst size and shape

The size of transmission bursts used in the experiments documented here was 36 slots. The smallest transmission burst size possible in WiMAX is 1 slot. Therefore the burst size used here is relatively large. A large burst is more likely to be subject to interference than a small burst, but a number of small bursts require more overhead information to be transmitted in the DL-MAP than a single large burst because each burst requires a separate entry. Hence, whilst it is probable that upto some point the relative performance gains obtained with frequency selective scheduling should increase as the block size is reduced, the increase in overhead may not compensate for the gains. It would thus be interesting to investigate the effect on performance of different burst sizes and find the optimal trade off between burst size and associated DL-MAP overhead.

It would also be interesting to study the relationship between burst shape, time diversity, and frequency diversity and Band AMC.

It is interesting here to note that in LTE [3], the size of a slot is fixed to 12 subcarriers x 7 OFDMA symbols (or 6 OFDMA symbols if the extended prefix is used), whereas a Band AMC slot in WiMAX is 9 subcarriers x 3 OFDMA symbols. Further, the maximum slot size for a single burst in LTE is 2 subframes, or 14 OFDMA symbols, whereas in WiMAX a burst can potentially span the whole DL subframe, which could be more than 14 OFDMA symbols. Thus, there are differences in the sizes and shapes of burst that LTE and WiMAX can transmit. These subtle differences could manifest as differences in performance for the same scenario, depending on how close to optimal burst size and shape each technology can realize.

• Channel estimation accuracy

Frequency selective scheduling depends on accurate channel estimation because it performs best when each user is given the radio resource that they will receive at the highest power. In the experiments performed here, channel estimation was considered to be ideal. This does not imply perfect reconstruction of transmitted symbols because there are noise sources such as thermal noise and inter-symbol interference which are non-invertible given channel estimation information. However it should be assumed that by employing ideal channel estimation, the results provide an upper bound on performance, since any degradation in channel estimation is likely to degrade performance. The effect of introducing imperfect channel estimation on the performance gains should be investigated.

• CQI feedback delay

Frequency selective scheduling uses the channel estimates of different resources obtained from transmissions that occur in the past, to influence the allocation of resources to users in the future.

Thus the channel state at the point of measurement may not be the same as the channel state at the point of transmission using that information. The larger this gap, the less likely it is that the measured information will still be valid. This gap is known as the feedback delay.

In this work the feedback delay was one frame, which is the theoretically smallest feedback delay that WiMAX can realize. It is assumed that the CQI estimates obtained from the DL portion of frame n are fed back in the UL portion of frame n and hence available to the scheduler for frame $n+1$.

It may not be practical to feedback CQI information every frame and feeding back CQI information consumes uplink resources, impacting the overall data throughput. Further, it may be redundant to feedback CQI information every frame if it is changing infrequently. Thus it seems prudent that further work should investigate the positive impact on performance of increasing the feedback delay. A scheme such as in [5] where the feedback delay is dynamically adjusted could be employed.

It is worth noting that LTE allows CQI reporting in principle once per subframe (every ms) compared to every 5ms in WiMAX, so that LTE has a greater flexibility in this regard and hence frequency selective scheduling maybe more efficient in LTE.

It is interesting to ask how these results might generalize to MIMO systems. Adding additional spatial dimensions, gives a scheduler more choices for burst selection. Consequently, a user with no high throughput bands in one spatial dimension, may gain such choices in multiple spatial dimensions. More choices for burst selection may also reduce the likelihood of ``decision collisions'', whereby two users with the same priority have the same preference for a band, forcing the scheduler to randomly downgrade one of them. Therefore future work should examine the notion of frequency-space selective scheduling.

More generally, increasing the bandwidth will increase the number of burst choices available to the scheduler. It would therefore be interesting to investigate whether frequency selective scheduling gains scale with bandwidth.

7 Conclusion

Frequency selective scheduling is a viable means of increasing throughput in WiMAX and throughput can almost be doubled at low velocity under ideal conditions. The work here demonstrates that a frequency selective scheduling scheme such as the WiMAX Band AMC mode of operation should be selected for urban cells in which the traffic source is dominated by pedestrians or stationary users. In rural environments, in which a system is deployed specifically to serve roads or railways, a frequency interleaving approach such as the PUSC mode of operation should be selected. It remains to be seen what effect more realistic constraints will have on the performance. The impact of other contributing factors and their interactions should be the subject of future work.

Chapter 4

The Impact of Scheduler Choices on the Performance of Static Intercell Interference Coordination Mechanisms

Group Research & Development

C2 – COMPANY CONFIDENTIAL

The Impact of Scheduler Choices on the Performance of Static Intercell Interference Coordination Mechanisms

Reference	VFGR&D_2010_0045
Version	1.0
Date	22 April 2010
Author	Ashley Mills
Theme/Function	Radio
Filed As	R&D Sharepoint https://sharepoint.vodafone.com/global/GroupRnD
Document type	Technical report
Status	Approved
Approved	David Lister
Signature / Date	.... / ....

"This page has intentionally been left blank (apart from this sentence, headers, footers, etc)."

Scope

This report addresses the choice of schedulers and their impact on interference mitigation techniques related to LTE.

3G/UMTS is out of scope.

Document History

Version	Date	Reason
1	22 April 2010	First Draft

Document Distribution

David Lister
Eric Murray
Neil Scully
Trevor Gill
John Turk
Peter Cosimini
Grahame Collins
Amy Lu

Executive Summary

This work examines the impact of changing cell scheduling strategy on the performance of static soft-reuse interference coordination schemes in LTE. It is found that the scheduling strategy chosen, and thus the operator policy applied, has a critical impact on whether or not a given static soft-reuse interference coordination scheme produces benefits.

The reason for this outcome is clearly explained in terms of the interaction between user distribution, cell SINR CDF, and chosen MCS Codeset.

Whilst in typical cell configurations a large proportion of UEs do have low SINRs, and consequently benefit from the application of soft-reuse, the gain per UE is only a fractional increase in bits per symbol. This is offset by larger reductions in bitrate, albeit over a smaller number of UEs, on the remaining regions of the cell. The net effect is a mean reduction in cell throughput.

The impact of this observation is that the performance of soft-reuse is critically dependent on the scheduling strategy and the control that scheduling commands over the relative proportional use of MCS schemes. In summary:

1. Scheduling strategies which give proportionately more resources to low SINR UEs, do benefit from soft-reuse schemes. In this work only reuse three was shown to benefit.
2. Scheduling strategies which give proportionally more resources to midrange and high SINR UEs do not benefit from soft reuse schemes, and thus obtain better cell capacities under reuse one.
3. For a fixed rate scheduler examined, the proportional use of MCS schemes was such that there were little practical differences between reuse one and the use of soft reuse schemes.

From these results it follows that the application of statically applied soft reuse schemes to large numbers of sites is a bad idea. There are too many realistic scenarios where net losses would be experienced, and too few scenarios which would obtain significant gains. **Soft reuse does not offer the fix-all solution that some vendors have formerly suggested, and thus it should not be generically applied.**

Instead, networks should be deployed in reuse one configurations unless it can be demonstrated empirically for particular markets that user distributions and traffic use patterns categorically favour the application of soft-reuse schemes. The only way to ascertain the latter is via direct analysis of the target network.

The recommendation advanced here therefore is to abandon the static soft-reuse approach in favour of research into dynamic inter-cell interference coordination schemes. These schemes will monitor for and identify the limited conditions under which coordination benefits can be obtained and react only as and when needed to coordinate the use of resources between affected UEs.

An approach along the lines proposed by Eric Murray in his report on *"An Analysis of a Load Balancing Handover Algorithm for LTE"* (VFGR&D_2010_0007), seems a promising direction to take. Load balancing can be viewed as a special case of interference coordination, and thus an integrated outcome may be feasible. In the proposed approach, resource orthogonality would be enforced between pairs of UEs in adjacent sectors only when this change has been identified as producing a net benefit.

Contents

1 INTRODUCTION

1.1 Motivation For Inter-Cell Interference Coordination

Shannon's seminal channel capacity theorem [23] dictates in the general case that the maximum supportable bitrate of a wireless channel having bandwidth B is exactly $B \cdot log_2(1+SINR)$. Where B is in Hz and SINR is the linear Signal to (Noise + Interference) ratio of the received signal.

For individual UEs, LTE allows the bandwidth B to be varied once every subframe, since each subframe is discretized across frequency and time into individually assignable units called Virtual Resource Blocks (VRBs). Therefore the throughput of a UE can be varied by modifying assigned B. B is of course upper limited by the availability of licensed spectrum, and in a multiuser system, fairness considerations dictate that individual UEs may not consistently consume the entire bandwidth.

For a fixed bandwidth assignment B it follows therefore that improvements in channel capacity must come from improvements in SINR. Of the factors that comprise the ratio SINR, noise is the least amenable to change. A so-called noise floor exists due to hardware cost considerations, and thus from a practical standpoint, noise is fixed.

Furthermore, in sparsely populated networks noise may be the dominant component of SINR, in which case the network is said to be noise-limited, however this is rarely the case in modern cellular networks. Thus it is the signal and interference components of SINR that draw focus.

In observing that interference arises precisely due to the signals of other transmitters, it follows that among a set of mutually interfering transmitters, a collective signal rise, will generally result in a collective interference rise. In such situations, increasing signal power has no effect on the SINR and the system is said to be interference-limited. This is the regime under which modern cellular networks operate.

For these reasons, the interference component of SINR has attracted a large amount of attention, specifically directed toward the goal of reducing it in the hope of improved user and cell throughputs.

Excluding MultiCell Processing (MCP) techniques, which attempt not to mitigate interference but to exploit it, every interference reduction technique is concerned with improving the orthogonality of received signals with respect to their interferers. Transmitted signals can interfere across multiple physical dimensions, and thus orthogonality appears in several different forms (see for example [7]):

- **Temporal orthogonality**

  Signals transmitted with sufficient temporal spacing do not interfere with each other.

  The temporal spacing qualifier is necessary here because wireless signals are subject to multipath reflections and thus can intefere with themselves and other transmissions over a longer time than the actual duration of the transmission.

- **Frequency orthgonality**

  Signals transmitted on non-overlapping frequency sub-bands do not interfere with each other.

Many interference coordination techniques hope to improve interference by coordinating the use of a shared resource so that its users are scheduled on different frequencies.

This is related to B in the Shannon equation. When B is under utilized in all cells, each cell may be given exclusive access to one or more different frequency sub-bands. But when B is fully utilized in one or more interfering cells, perfect frequency orthgonality can only come at the cost of reduced B in those cells.

In this sense then, frequency orthogonality and B are often the subjects of a tradeoff.

- **Spatial orthogonality**

Signals from sufficiently separated transmitters do not interfere with each other.

This claim follows from fundamental physical laws. Signal power decays exponentially with distance, and so with sufficient separation, signals from different transmitters do not interfere with each other. In practice, signals are subject to shadowing and other effects, but on a basic level the claim holds.

This obvious claim motivates the use of power control in interference coordination. For example, since lower power transmissions travel less distance than higher power transmissions, soft frequency reuse, discussed in detail later, uses low power transmissions to spatially separate the centres of interfering cells, that share the same transmission bandwidth. Spatial separation also appears in uplink power control.

Spatial separation is often achieved via antenna directionality. For example, each sector in multi-sector site has its own antenna whose parameters are tuned to minimize coverage overlap with adjacent sectors.

At a more sophisticated level, antenna directionality can be changed on a per transmission basis and at a much higher spatial resolution. This is called beam forming and is an interference coordination approach that has become popular recently due to contemporary technical advancements.

- **Code orthgonality**

The dot product of orthogonal vectors is zero, the dot product of a vector with itself is one.

Although not relevant to LTE, except in the encoding of pilot symbols, code orthgonality forms the basis of the UMTS cellular systems. UMTS systems deliberately occupy the same frequency, time, and spatial resources and yet use semi-orthogonal data encodings to ensure appropriate separation.

A final kind of orthogonality occurs at the level of the channel. MIMO receiver schemes employ multiple receive antennas precisely because the received signals are partially uncorrelated due to having taken different spatial paths. By combining antenna spatial separation with code orthogonality, elegant schemes such as those proposed by Alamouti [2] have been extended [17] to provide very effective means of improving UE bitrates.

Signal orthogonality comes at a price. At a fundamental level, base stations cannot be spatially separated by too much, otherwise handover would be impossible, and coverage holes would appear. Frequency orthogonality comes at the cost of reduced bandwidth B. Spatial separation through power reduction comes at the cost of reduced signal and thus throughput. Temporal orthogonality of interferers results in a reduction in the availability of viable timeslots for transmission, and thus in effect, bandwidth again.

The general price for reduced interference, is usually therefore reduced resource. Even in the case of advanced beam forming separation techniques and MIMO schemes, the cost can be seen in increased computational complexity and real monetary costs of hardware.

The central question that must be addressed in all these cases, is whether or not the benefits of resource orthogonality are worth the resource cost.

A classical example of this question is whether or not it is better to deploy a system in a reuse three configuration or in a reuse one configuration. Under reuse three, the SINR of each cell is vastly improved compared to reuse one due to frequency orthogonality between otherwise interfering cells, yet these SINR improvements come at the cost of having to divide the available bandwidth into three orthogonal sets and thus a factor three reduction in B. Which scenario achieves higher throughput depends on the relative increase due to improvements in SINR which the imperfect orthogonality brings, compared to the loss in throughput due to the reduction in B.

It should be clear that if B is reduced by a factor of three, then the throughput due to improved SINR must be increased by more than a factor of three for it to be a worthwhile enterprise [1].

The compensation that improved SINR brings when frequency orthogonality is enforced at the behest of B, is so named the ``compensation effect'' in [19]. But conversely, one could say that reduced SINR can be compensated by increasing B. And so interference coordination is more about tradeoffs and choosing between multiple available options, than the one-way compensation of one thing for another.

In the sections that follow, this work will examine the performance of a particular type of interference coordination called soft-reuse, otherwise known as soft frequency reuse or fractional frequency reuse.

1.2 Summary of existing work and contribution of this work

The majority of static soft and fractional reuse schemes proposed to date (see for example [20, 21, 8, 26, 10, 11, 24]) are based on a similar premise. The schemes are based on the idea that UEs at the cell's centre can operate under a reuse one scheme because their spatial separation results in little mutual interference, and that UEs at the cell's edge however would benefit from operating under a tessellated reuse scheme, such as reuse three, since they are more impacted by the mutual interference generated by adjacent cells. This premise is echoed in the outcome of simulations, and has resulted in strong claims being made by vendors as to the benefits of soft reuse. For example, in [8] the authors claim that the application of soft reuse results in "Improved bit rate at cell edge" and "High bit rate at the cell centre". The implication is that soft-reuse is somehow able to get the benefits of both reuse one and frequency coordination.

[1] Not considering other restrictions such as minimum required service SINRs, as in GSM.

Unfortunately the body of work in general suffers from a lack of statistical rigor, and generally falls short of stating all important assumptions. In the case of 3GPP work, the situation is made worse because submissions are not subjected to the same degree of peer review as are academic papers.

Furthermore, when benefits are expounded, there is little explanation as to why or under what restrictions the claimed benefits apply. This makes it very difficult to interpret the body as a whole and make practical recommendations that are relevant. In addition, most of the work is done in hexagons.

To give an example, consider [19], this paper is titled to suggest that its technical assumptions are aligned with 3GPP LTE standardization yet it uses a non-standard definition of resource block. Furthermore, and more importantly, whilst the paper claims to compare interference coordination with interference randomization, it really only demonstrates the somewhat obvious outcomes that (i) cells will experience interference if their resources are not frequency-orthogonal, and (ii) frequency diversity gains scale with bandwidth. This results in a tiny diversity gain for interference randomization at loads which guarantee non-orthogonality, that amounts to between a ¼ and ½ dB decoder gain at BER rates below typical MCS reselection targets. It is impossible to tell how this maps to system throughputs since only two extremely low MCS rates are studied and the system setup consists of only 4 cells.

This work mitigates the problems identified by providing a number of things:

1. Statistical rigor. All qualitative comparisons are quantified in terms of appropriate statistical tests.

2. Realistic scenario. Results for a central London scenario are provided alongside hexagonal results.

3. Explanation. Soft reuse is not only shown to demonstrate a gain, it is explained exactly why and under what possible circumstances it can demonstrate a gain.

1.3 Document outline

The next section describes the procedure and results for a number of experiments using soft-reuse. The experiments have been designed with the goal of understanding under what circumstances soft reuse can provide a benefit.

Section 2.1 describes in detail the underlying simulation assumptions. The general idea is that a number of soft reuse power profiles will be examined for different scheduling strategies and other cell modifications.

Section 2.2 looks at how the application of different soft reuse power profiles, changes the mean cell performance. That is, the performance under a given soft reuse power profile averaged over the whole area of the cell.

2.3 extends the results to look at the ability of different soft reuse power profiles to support fixed rate UEs. That is, a number of UEs all with the same fixed rate target.

2.4 generalizes the results by exploring for a multiuser scenario, what happens when the resource allocation is biased toward UEs in particular parts of the SINR CDF. This is framed in terms of throughput fairness.

2.5 discusses the results and 2.6 concludes.

2 EXPERIMENTS IN SOFT REUSE

2.1 Simulation Methodology

2.1.1 Network Layout

2.1.1.1 Hexagonal dataset

In the interests of making the presented results accessible and easily replicable, and in accord with common practice, simulations have been performed for a classic simulated hexagonal deployment.

A two-tier hexagonal network containing 57 cells with wraparound interference was used as illustrated in Figure 1

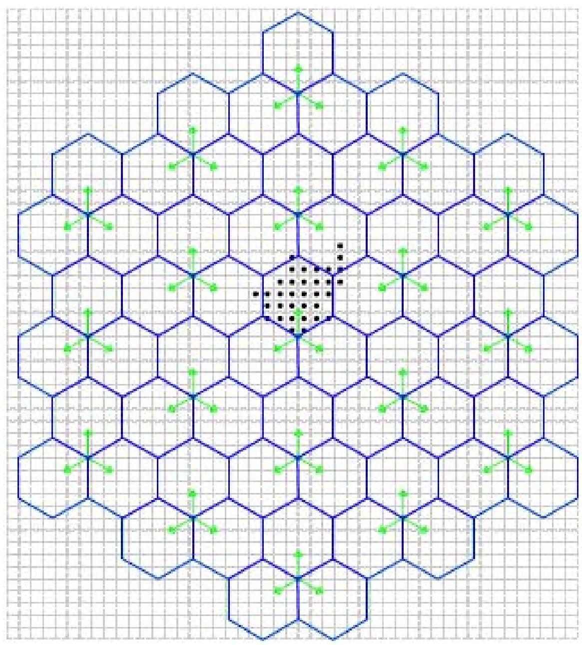

Figure 1: A 57 cell hexagonal deployment shown with a grid resolution of 50m2 for clarity. Only the centre cell is populated with UEs. In the experiments a resolution of 5m2 was used

The centre cell of this deployment was singled out for focus and was populated with UEs at a density of one UE per 5m^2 over its service area. The cell's service area was defined as the area over which received signal from that cell (RSRP) was stronger than from any other cell. RSRP was computed at a resolution of 5m^2 such that a sampling of all the UEs in cell was equivalent to sampling the whole service area. All cells that were not the focus cell were set to transmit at full power continuously, restricted only by the given soft-reuse power profile under study, in order to represent the worst case scenario for interference coordination. The parameters shown in Table-1 were used to establish the hexagonal environment.

Parameter		Value	Unit
Number of Cell Sites		57	
Number of Sectors Per Site		3	
Inter-Site Distance		500	metres
Lognormal Shadowing Standard Deviation		7	dB
Shadowing Correlation	Between Cells	0.5	
	Between Sectors	1.0	
BS Antenna Boresight Gain		14	dBi
BS Horizontal Antenna Pattern		$A(\theta) = -\min\left[12\left(\frac{\theta}{\theta_{3dB}}\right)^2, A_m\right]$ θ_{3dB} = 70 degrees, A_m = 20 dB	dB
Carrier Frequency		2000	MHz
System Bandwidth		600 subcarriers x 15 kHz = 9000	kHz
Distance-Dependent Path Loss		COST231 Extended Okumura-Hata Model (URBAN clutter scenario) [1]	
Base-Station Antenna Height		30	metres
UE Antenna Height		1.5	metres
UE Noise Figure		9	dB
Minimum Coupling Loss		60	dB

Table 1: System parameters for the hexagonal deployment scenario

During experiments, data was collected from every UE within the focus cell.

To guard against statistical aberration, all hexagonal results presented were averaged over 40 random shadow fading instances. Where one instance corresponds to one realization of the 57 cell grid.

2.1.1.2 London dataset

In order to mitigate the inevitable criticism levied toward hexagonal datasets, data for a realistic London scenario has been obtained so as to serve as a comparison to the hexagonal case.

The data represents an area of central London close to Paddington Station. Antenna settings and terrain data reflect the actual network settings as of 2004 for the Vodafone UMTS deployment.

Pathloss was calculated at a resolution of $10m^2$ using the Pace3D ray tracing software module in Atoll [4]. Pace3D accurately models the effects of building penetration losses, reflection, and refraction effects and provides a realistic picture of the actual pathloss variation experienced in each cell.

The exact area is shown in Figure 2. The top left of this area has a Latitude,Longitude coordinate of 51.0817871,-0.139529675, and the bottom right are has a Latitude,Longitude coordinate of 51.53060532,0.193410516.

Figure 2: Hybrid Road/Satellite map of London area used for simulations

Figure 3 shows the Atoll generated pathloss predictions for this area.

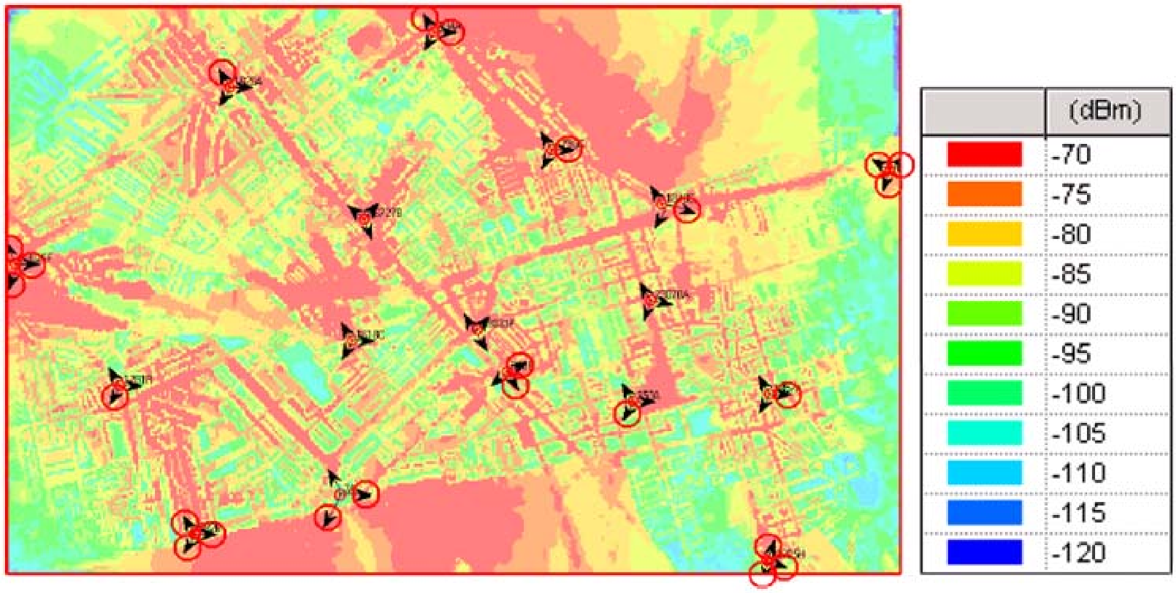

Figure 3: Atoll predictions and antenna locations for the London area shown in Figure 2. Circled antennas are excluded from results collection

Since inter-site interference was generated only for those sites pictured, it would be unrealistic to use the sites situated at the edges. Because these sites receive little inter-site interference, they are not representative, and including them could bias the outcome of experiments. To address this problem a number of sites have been conservatively excluded from averaging in experiments and serve only to generate interference to the other sites. The circled antennas shown in Figure 3 are those which have been excluded. All remaining sites are averaged over during experiments, to protect against statistical aberration.

2.2.2 MCS Codeset

An MCS codeset is a set of modulation and coding schemes to be used under link-adaptation. A set of MCS schemes is needed because each MCS is optimal for a particular narrow SINR range. Two MCS codesets are employed in this work. The first codeset is taken from Appendix A2 of a 3GPP LTE standards simulation methodology document [14], it is thus referred to as the 3GPP codeset in this document. The 3GPP codeset is shown in Table-2.

MCS	Rate (bits per symbol)
QPSK 1/8	0.2500
QPSK 1/5	0.4000
QPSK 1/4	0.5000

QPSK 1/3	0.6667
QPSK 1/2	1.0000
QPSK 2/3	1.3333
QPSK 4/5	1.6000
16QAM 1/2	2.0000
16QAM 2/3	2.6667
16QAM 4/5	3.2000
64QAM 2/3	4.0000
64QAM 3/4	4.5000
64QAM 4/5	4.8000

Table 2: 3GPP Reference Codeset

For system simulations any given MCS codeset needs to be associated with a lookup curve which maps SINR values to MCS schemes, and thus throughputs. The lookup curve used here is also derived from [14]. The curve is based on an AWGN channel and shannon throughput, that has been reduced by a factor $3/4$ and then quantized to the above codeset. The mapping is shown in Figure 4.

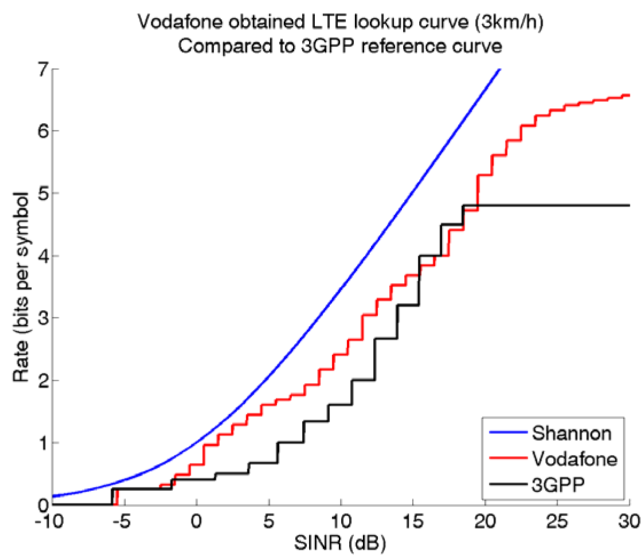

Figure 4: Vodafone obtained SINR to MCS lookup curve compared to a 3GPP reference curve

In addition to the 3GPP lookup curve, a different a lookup curve is also examined. The motivation for this is explained in Section-2.3 but it makes for a more concise document to introduce the lookup curve here. The mapping is shown in Figure 4 against the 3GPP curve. Since the curve is derived from a Vodafone source, it is referred to as the VF curve in this document.

This curve is obtained from Alcatel-Lucent and is based on a Pedestrian B channel and assumes 3km/h mobility. It is considered to be an accurate representation of the throughputs that the real LTE hardware will deliver.

The first major difference with the 3GPP reference curve is that MIMO 2x2 is used, hence an improvement in rates across the majority of the SINR range is seen. The second major difference is that the 3GPP reference curve does not offer a rate improvement beyond 18.5dB whereas the Vodafone curve offers rate improvements up to 30dB.

There is no real codeset associated with this lookup curve since it is created from a number of aggregated transmissions. However, the quantized function used is provided in Table 6 of Appendix A.1 of this document to facilitate external replication of results.

2.1.2 LTE System Assumptions

Figure 5 illustrates the essential components of a DL LTE frame (in 10Mhz). The frame itself consists of 10 subframes which each last 1ms. Half of a subframe is called a slot. In frequency, each subframe is split into 50 Virtual Resource Blocks (VRBs). To the right of the frame a VRB is shown expanded.

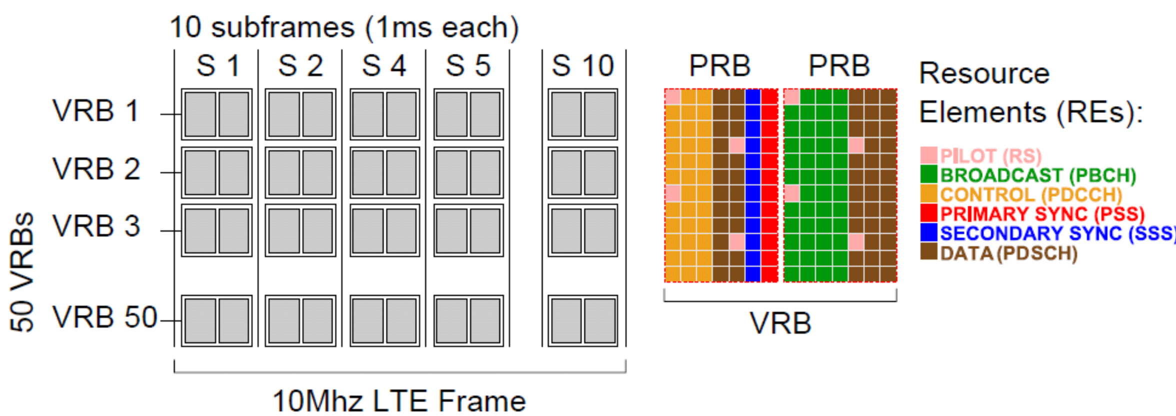

Figure 5: The essential components of a DL LTE frame (10Mhz Bandwidth)

Each VRB is comprised of a pair of Physical Resource Blocks (PRBs). Section 6.2.3 of 3GPP 36.211 [13] states that PRBs are mapped to VRBs under a distributed or localized mapping, and are assigned to UEs in pairs (one from each slot of the subframe) under a single VRB number. Accordingly this text uses the VRB as an atomic unit of allocation, since it is in fact technically impossible to independently allocate the PRBs within a VRB. Each PRB of a VRB spans 12 subcarriers in frequency and 7 symbols in time (note that for large inter-site distances a longer cyclic prefix is available which spans only 6 symbols in time, but that here the former duration is used). Each unit in a PRB is called a Resource Element (RE). An RE spans a single subcarrier in frequency and single symbol in time. An RE has a frequency width of 15khz and lasts approximately 70µs.

An RE encodes exactly one MCS symbol. For this reason the term symbol is often used interchangeably with RE. LTE uses REs for a number of different purposes in addition to serving data. The RE types modelled here are shown, with their locations in the VRB to the right of Figure-5. The following list describes their functions.

- Control channels

Control channels are used here only to model VRB RE usage. Decoding of control channels and associated decoding failures are not modelled. Control channels are assumed to take up the first three RE columns of each VRB.

- Pilot symbols

Pilot symbol placement is modelled accurately according to [13]. The computation of RSRP from pilot symbols is performed as defined in Section 5.11 of [15]. RSRP measurements are used to associate UEs with serving BSs in the static simulations. The figure shows the pilot symbol positions for a single spatial channel on a cell having CELL_ID = 0.

- Synchronization and broadcast channels.

Synchronization and broadcast channels are included only to model VRB RE usage accurately. Primary and secondary synchronization symbols are only present on the 3 VRB either side of the central carrier and are only present in subframes 0 and 5. Broadcast channels are only present on the 3 VRB either side of the central carrier and are only present in subframe 0. They are shown in the VRB diagram only to indicate their positions when they do occur, but be aware that most of the time data REs take their place.

As already stated, synchronization and broadcast channel mechanisms are not simulated, they only consume space, so as to model the RE usage accurately.

- Data symbols

Since only subframes 0 and 5 contain synchronization and broadcast channels, the number of data sresource elements per VRB, changes depending on the subframe and location within the band. Taking this into consideration, an average number of data REs per VRB can be computed.

Factoring in the presence of control and pilot symbols, the average number of data resource elements per VRB for the simulator is 124.8720 data symbols. This is the figure which is at the root of all throughput calculations.

As an example, consider the MCS 64QAM-3/4, this MCS has a rate of 4.5 bits per symbol. Therefore a single VRB on average at this rate can carry 561.9240 bits. This corresponds to a peak rate of $\approx 28Mbps$.

2.1.3 Definition of Soft Frequency Reuse

Unlike LTE, which is standardized, there is no common concensus as to the definition of soft-reuse. Furthermore, soft reuse is sometimes presented [6, 27, 8, 20] in the misleading manner shown in Figure 6.

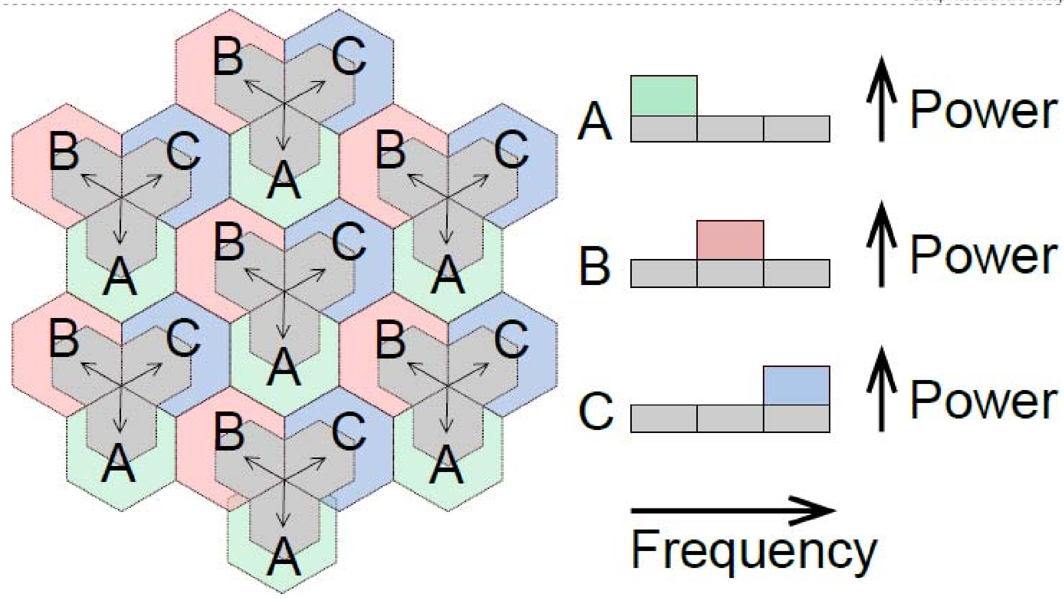

Figure 6: Soft reuse as conventionally presented

The left side of the figure shows a hexagonal configuration of cells, colour coded according to a tesselating pattern indicating which parts of the frequency band are allowed in each part of each cell. The right of the figure shows the frequency-power transmit profiles for each of the three types of cell that arise. The general idea is that:

a. UEs close to the BS can use all of the band to transmit, albeit at a lower TX power. A lower TX power is sufficient because the UEs are close to the BS. The whole band can be used because the low power transmission is unlikely to interfere with other cell-centres.

b. UEs far from the BS use a designated one third of the band, which is coordinated in a tesselating pattern with neighbor cells in order to reduce interference. The frequency orthogonality with neighbors improves SINR and throughput for the otherwise low-SINR cell edge UEs and allows a higher transmit power to be used, which boosts cell-edge throughput still further.

This conventional presentation is misleading because it suggests a particular correspondance between the geographical position of a UE and the spectrum it should use. However, there is no evidence to suggest that this particular assignment of UEs to spectrum is optimal for all possible performance metrics, and indeed the no free lunch theorem [12] makes it seem likely that there cannot be. With this in mind, in this work soft-reuse has been conceptualized in a manner which is free from implying a particular correspondance between UEs and spectrum. Figure 7 illustrates the concept.

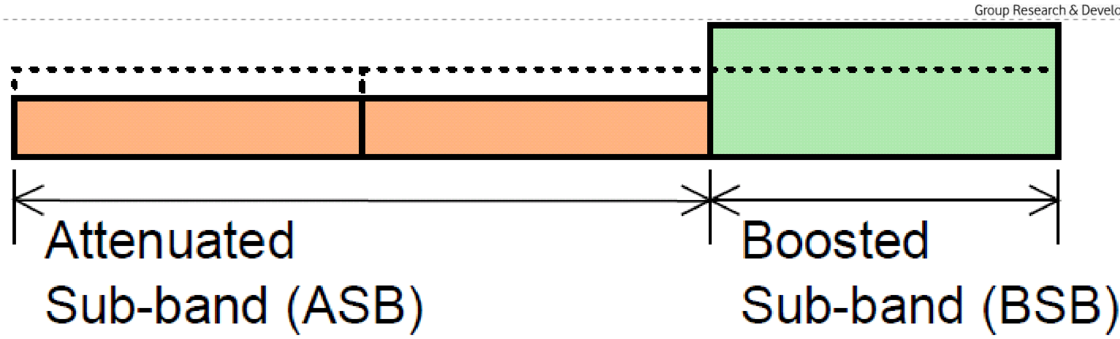

Figure 7: Soft Frequency Reuse Configuration

In this case soft reuse is simply described as having part of the spectrum boosted and part of the spectrum attenuated (where the special case of the spectrum being entirely attenuated is allowed).

The band is split into two unequal parts in the ratio 2:1, this ratio was chosen because: firstly, it allows coordination in a manner comparable to reuse three. And secondly, it is a commonly used reference case, see for example [25, 7].

These two bands have been named so that each region can be referred to as a noun. Thus the attenuated part of spectrum is given the name Attenuated Sub-band and referred to by the mnemonic ASB, whereas the boosted sub-band takes the name BSB.

The ASB can be used by all cells and the BSB is setup with a reuse pattern of three among cells. By changing the ratio of power on the ASB and BSB, various ``soft-reuse'' schemes can be established for example:

- Reuse one. The power per VRB on the ASB and BSB are set equal.
- Reuse three. The power per VRB on the ASB is set to zero, and the power per VRB on the BSB is set to three times the reuse one value.
- Soft reuse. The power per VRB on the ASB is reduced, but not to zero, and the BSB is boosted proportionately to the power taken away from the ASB.

2.1.4 Soft Reuse Power Ratios Examined

Reuse three can be thought of as lying at one end of the soft reuse spectrum, where the ASB power is set to zero, and reuse one can be thought of being at the other end, with ASB and BSB powers being equal. To investigate the range of soft reuse power ratios possible, the power ratios shown in Table-4 were examined.

Reference Index	1	2	3	4	5	6	7	8	9	10
Power per VRB (ASB)	0	1/9	2/9	3/9	4/9	5/9	6/9	7/9	8/9	1

Power per VRB (BSB)	3	25/9	23/9	21/9	19/9	17/9	15/9	13/9	11/9	1

Table 3: The soft reuse power ratios used in the experiments

These power ratios are depicted graphically in Figure 8.

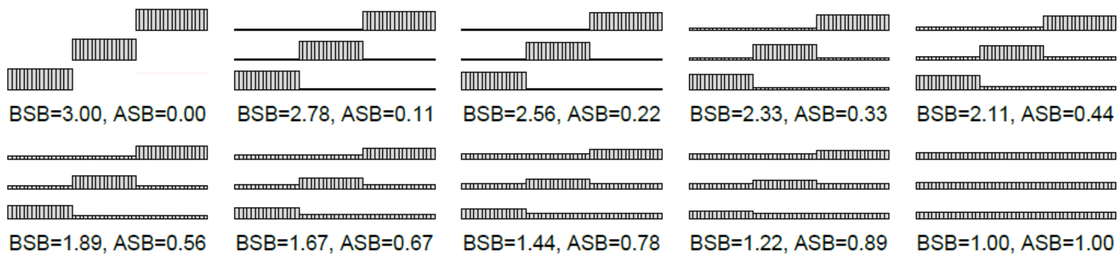

Figure 8: The soft reuse power ratios used in the experiments

Note that all of these power ratios consume a total of 40W per sector.

2.1.5 Summary of Simulation Methodology

All experiments in this chapter consist of examining the behaviour of a number of soft reuse power ratios applied to a simulated LTE cellular network for one hexagonal and one London dataset. In the sections that follow, the impact of changing the soft reuse power ratio will be examined with respect to various different scheduling strategies, so as to gain an understanding of under what circumstances, if any, soft reuse obtains a benefit.

2.2 Mean Bitrate Scheduling

2.2.1 Motivation

This section examines how the interference conditions under different reuse power ratios change the mean cell bitrate. The mean cell bitrate is calculated from the average bit rate over all grid positions within the coverage area of that cell.

2.2.2 Procedure

The procedure used for obtaining the data is outlined in Procedure 1.

```
ForEach (Soft Reuse Power Ratio) : Do
    For (Shadowing Seed) In (1 to 40) : Do
        ForEach (Focus Cell UE) : Do
            A <— ASB Bitrate
            B <— BSB Bitrate
            C <— Mean Bitrate, 2/3 * A + 1/3 * B
        Done
        Compute mean of A,B,C over all Focus Cell UEs
    Done
    Compute mean of A,B,C over all Shadowing Seeds
    Record A,B,C for this Soft Reuse Power Ratio
Done
```

Procedure 1: Procedure used to generate mean cell throughput results

The procedure shown states that 40 shadowing seeds were averaged over. This applies only to the hexagonal dataset. For the London dataset there is only one shadowing seed so to speak; i.e the actual shadowing distribution as predicted by Atoll. Thus for the London dataset, results were averaged over all non-excluded focus cells in the scenario, as explained in Section 2.1.1.2.

2.2.3 Results

Figure 9**a** shows the mean cell throughput obtained under each soft reuse power ratio for the hexagonal dataset. The mean is as averaged over the service areas of 40 shadow fading instances, as outlined in Procedure 1. Figure 9**b** shows the mean cell throughput broken down into its ASB and BSB components.

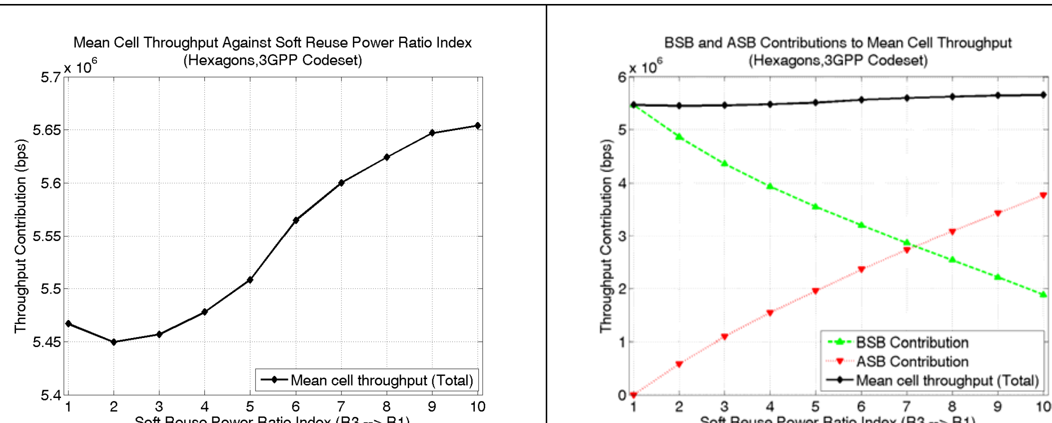

(a) Mean cell throughput. **(b)** ASB/BSB contributions.

Figure 9: (a) Mean cell throughput under different soft reuse power ratios. (b) The relative contributions to the mean cell throughput from the ASB and BSB. (HEXAGONAL data, 3GPP MCS codeset)

Figure 9**a** shows that the highest mean cell throughput is obtained under reuse one and shows a general decline in mean cell throughput as the power ratio is changed in the direction of reuse three. The mean cell throughput under reuse one is 3% larger than under reuse three.

To test for the significance of this result, a one tailed t-test was performed. Since it cannot be assumed the variances of the input samples are equal, Satterthwaite's approximation was used to address the Behrens-Fischer problem. The outcome of this test had a *p* value of 0.16, which means that there is a 16% chance that the differences observed are due to chance, under the assumption that the results are normally distributed. Instructions on how to perform such a test using the Matlab statistics toolbox are provided in Appendix D.

All quantitative comparisons which follow in this section will undergo the same test for significance, and in each case the *p* value will be provided. A *p* value less than 0.05 is often used as a threshold to reject the null hypothesis, but it should be remembered that a *p* value only provides a degree of confidence to the reader which they should judge accordingly themselves in light of the other presented assumptions [5].

For example, the improvement of reuse one over reuse three shown in Figure 9 was 3% with a p value of 0.16. Although this value is not less than 0.05, when the overall trend of the graph is considered, it is reasonable to believe that there really is a small difference between reuse one and reuse three. Whereas if the p value had been something like 0.5, it would be too high to make such an inteference.

The relative contributions from the ASB and BSB that make up the mean cell throughput are shown in Figure 9**b**. The observed decline in mean cell throughput as the interference moves from reuse one towards reuse three along the abscissa comes about because: the rate at which the BSB throughput improves due to improved power per VRB on the BSB does not compensate for the rate at which throughput declines on the ASB due to reduced power per VRB on the ASB.

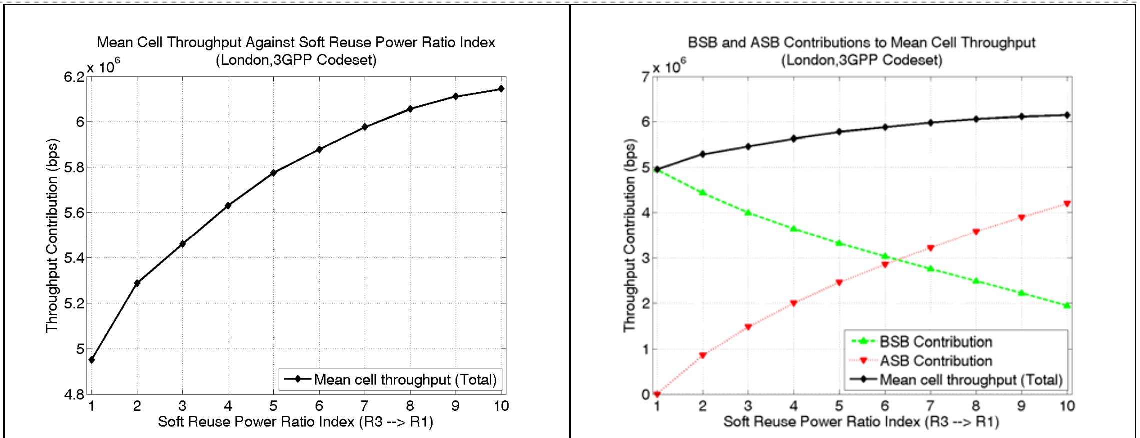

(a) Mean cell throughput. **(b)** ASB/BSB contributions.

Figure 10: : (a) Mean cell throughput under different soft reuse power ratios. (b) The relative contributions to the mean from the ASB and BSB. (LONDON Data, 3GPP Codeset)

Figure 10 shows the same plots, for the London data case. It is shown that the London case is consistent in behavior with the hexagonal case, in demonstrating the same trends. However, for the London case, the drop in mean cell throughput due to employing soft reuse is much greater than for the hexagonal dataset. Here reuse one improves on reuse three by $\approx 24\% \ \ (p \approx 0.001)$.

It is interesting to ask in this simple case, exactly why the mean cell throughput drops in the observed manner when moving away from reuse one, and is not ``compensated'' by improved frequency orthogonality.

Before adducing the reason for this, it will help to develop some intuition. Consider the difference between any cell under a reuse one interference profile, and the same cell under a reuse three interference profile. The cells differ in only two ways:

1. The SINR distributions are different.

2. The reuse one case has three times the bandwidth of the reuse three case.

Since SINR is quantized into acheivable rates by means of a lookup curve, this means that the area of a cell, in being covered by a particular SINR distribution, is also in fact partitioned discretely according to different supportable MCS rates. This is illustrated in Figure 11.

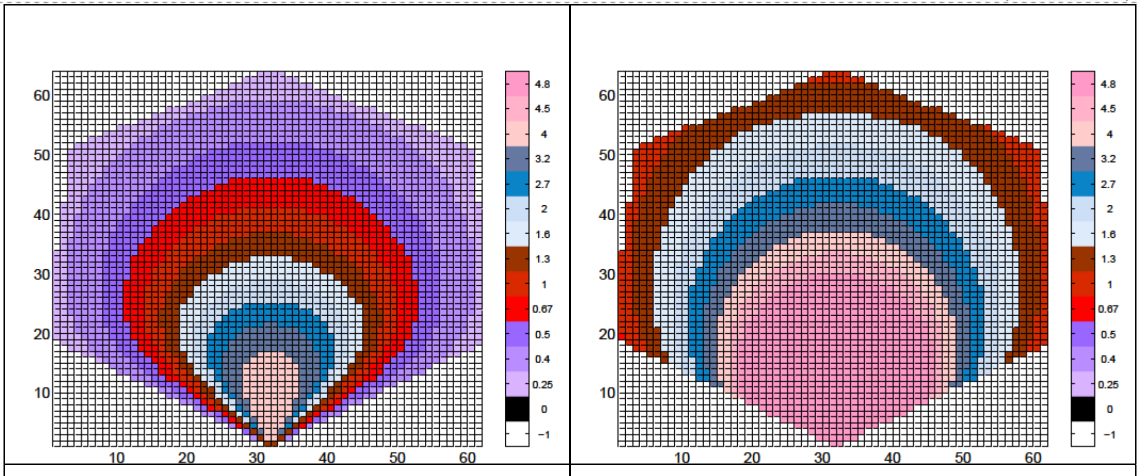

(a) MCS coverage under reuse one. **(b)** MCS coverage under reuse three.

Figure 11: Cell coverage shown as maximum supported MCS rate for reuse one and three. The colouring represents the maximum achievable MCS and the corresponding rates are shown in the colour scale to the right of each plot. Plot (a) shows reuse one and plot (b) shows reuse three.

Figure 11 shows the coverage, according to maximum acheivable bitrate under reuse one and reuse three scenarios, for a hexagon cell with no shadow fading. Note that these are the only plots in this section not to employ shadow fading, and are merely included for instructive purposes. From the figures, several observations can be made:

- SINR declines as a function of distance from the BS.
- Area increases as a function of distance from the BS.
- The areas served by lower order MCS schemes are therefore typically larger than the areas served by higher order MCS schemes.
- Reuse three benefits more from higher SINR than reuse one.
- When the SINR is improved under reuse three, for any given MCS coverage area under reuse one, the improvement will likely result in higher order MCS schemes being made available.
- The improvement in achievable MCS rate that occurs in a given MCS coverage area under reuse one, may vary across that coverage area.

In seeing these figures, it should be clear that the overall improvement to a cell that occurs when switching to reuse three, depends in a complicated manner on the interplay between the MCS quantization and how this maps to coverage across the cell in each case. The next figure, Figure 12 illustrates how these factors interact to produce the mean throughput plots of Figure 9 and Figure 10.

Remembering that each MCS rate is supported over a particular cell area under a given reuse scheme as illustrated in Figure 11, consider what happens to the acheivable rate within a given MCS area when the reuse scheme is changed from reuse one to reuse three. In general, since reuse three improves the SINR, it will be the case that over the same area, the acheivable rates will increase.

Figure 12 qualifies this statement for the hexagonal data. For each MCS area under reuse one, the average rate under reuse three was computed. Figure 12 thus plots the mean rate under reuse three for each MCS used under reuse one. The results are averaged over the cell areas of all shadowing seeds. The mean rate under reuse three has been reduced by 1/3 to reflect the reduction in bandwidth which occurs under reuse three, and is thus called the effective rate.

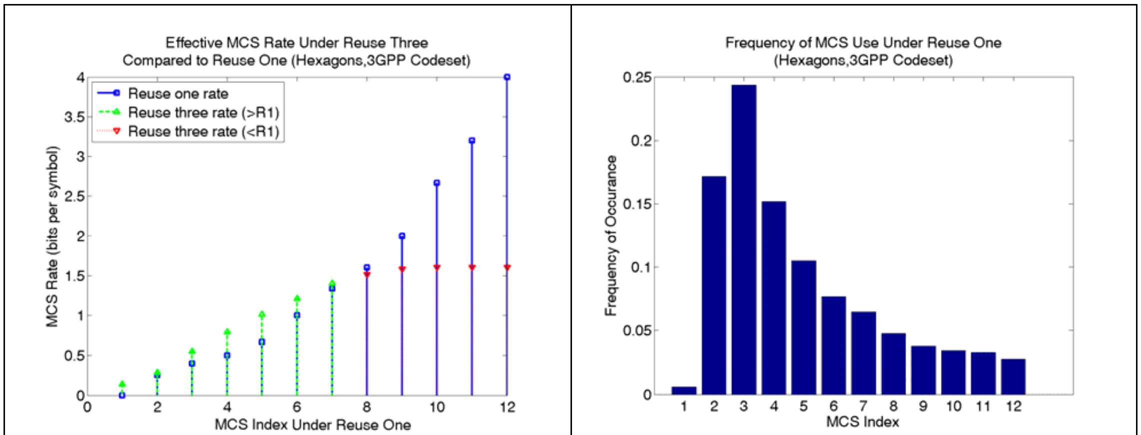

(a) Effective rate change when switching from reuse one to reuse three.

(b) Relative proportional use of MCS schemes by UEs.

Figure 12: (a) Effective change in obtainable bitrate when switching from a reuse one interference scenario to a reuse three interference scenario. (b) Relative proportional use of MCS schemes across the cell (Hexagons, 3GPP Codeset).

Figure 12**a** shows that for some parts of the cell, namely those supported by low MCS rates under reuse one, switching to reuse three incurs a benefit. These points are shown as up-arrows. For other parts of the cell, switching to reuse three results in a loss in throughput. These points are shown as down-arrows.

To make this explicit, consider the bar at MCS index 8. The blue squared line represents the bitrate of UEs that get MCS index 8 under reuse one. This of course corresponds exactly to the bitrate of index 8. Taking these same UEs, and changing the interference conditions to reuse three, the red down arrow represents the mean rate of those UEs under the new interference conditions, after a 1/3 bandwidth reduction has been factored. In this case the effective rate under reuse three is less than the rate under reuse one and so a loss is incurred for those UEs.

Whether or not reuse three beats reuse one overall consequently depends on the relative proportional use of each MCS. Figure 12**b** shows the frequency of use of each MCS under reuse one.

Computing the proportions, it turns out that over $\approx 81\%$ of the cell, reuse three beats reuse one, in that the resultant improved SINR supports an MCS with at least three times the rate under reuse one. However, the mean improvement in rate is only 0.17 bits per symbol. Over the remaining $\approx 19\%$ of the cell area, reuse three gets less than three times the throughput of reuse one, and suffers a mean reduction in rate of 1.12 bits per symbol. So while reuse three improves things for the majority of the cell, the improvement per UE is small in comparison to the loss incurred per UE on the remainder of the cell. Weighting each case arrives at the conclusion which supports Figure 9, namely that the overall winner in the mean case is reuse one.

The tradeoff is more apparent for the London dataset, and is shown in Figure 13.

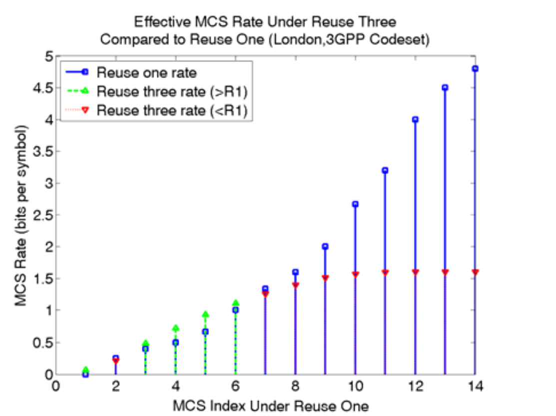

(a) Effective rate change when switching from reuse one to reuse three.

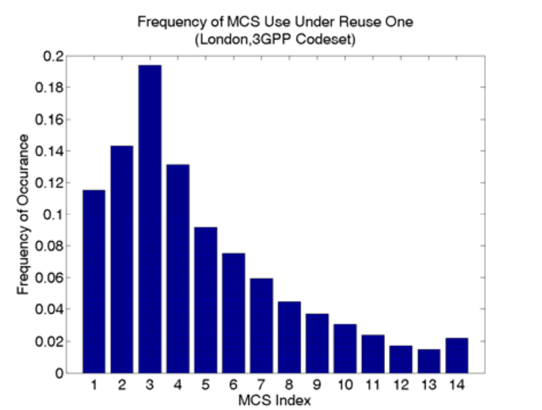

(b) Relative proportional use of MCS schemes by UEs.

Figure 13: (a) Effective change in obtainable bitrate when switching from a reuse one interference scenario to a reuse three interference scenario. (b) Relative proportional use of MCS schemes across the cell (London, 3GPP Codeset).

Here reuse three beats reuse one over $\approx 61\%$ of the area of the cell, with an average gain in rate of about 0.15 bits per symbol. Reuse one beats reuse three in the remaining 39% of the cell and the gain in rate is about 1.33 bits per symbol. The overall effect is that reuse one wins out.

It is interesting to note that in both these cases reuse three was able to compensate for the loss in bandwidth only when the SINR was low. This seems rational, because for low rates there is plenty of scope for improvement when switching to reuse three, yet at high rates, since the MCS curve used stops at approximately 18.5dB, the benefits of improved SINR stop.

This can be seen by looking at the empirical CDFs in both cases, as illustrated in Figure 14 for both the hexagonal and london datasets.

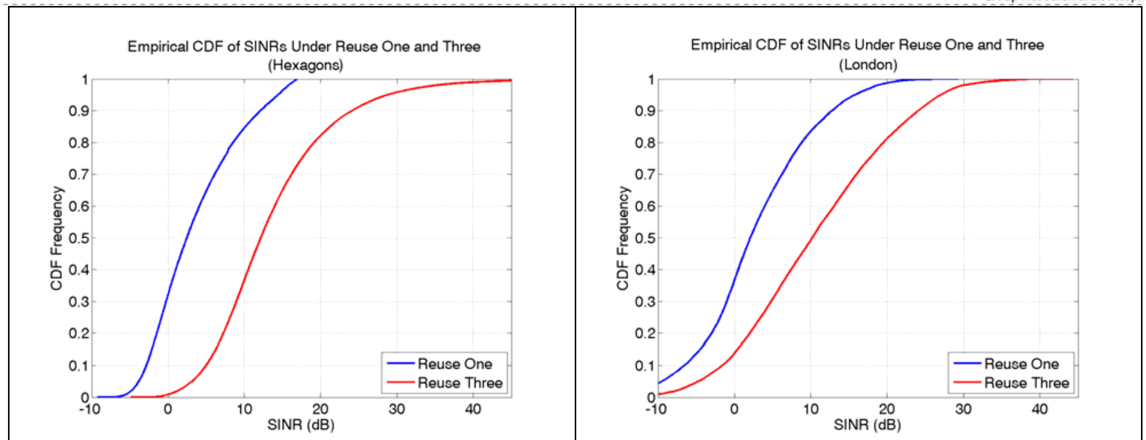

(a) Hexagonal SINR CDFs. **(b)** London SINR CDFs.

Figure 14: CDFs of cell SINRs under reuse one and three interference conditions.

Note, under reuse one in the London case, about 10% of the UEs have SINRs lower than -6.5dB, which is the minimum threshold for operation under the 3GPP MCS codeset. This is probably artificially high in this work for two reasons.

Firstly, in this work traffic was sampled uniformly from the London area whereas in reality the traffic distribution is unlikely to be non uniform. Since antenna tilts are optimized for the areas where traffic is served, it is likely that in the real network the CDF curves will be right shifted slightly.

Secondly, in real networks which are rarely saturated entirely, the average SINR of UEs will be improved and thus the curves will probably be right shifted. In addtion, time and frequency dependent variations will mean that a given UE will have some chance of being raised above the threshold on a per frame basis, even if their mean SINR is bad.

Looking at the top end of the distributions, taking the hexagonal dataset first, 97% of UEs have SINRs less than 18.5dB, which means that throughput is almost never limited by MCS codeset. Under reuse three however, 22% of the UEs have SINRs greater than 18.5dB and 18% have SINRs greater than 20dB. These UEs are likely to be limited by the MCS codeset since beyond 18.5dB it provides no further gains in throughput. This means that quite a large percentage of UEs would benefit if the MCS curve supported higher order MCS schemes, and suggests that the results may be different for a different MCS codeset.

In order to test this, an alternative MCS curve was obtained from an LTE equipment vendor. This curve was introduced earlier in Figure-4.

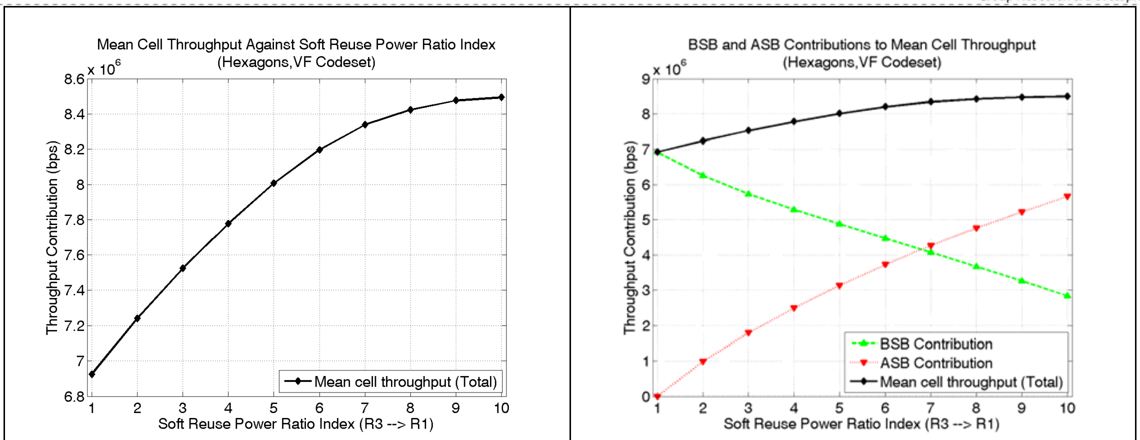

(a) Mean cell throughput. **(b)** ASB/BSB contributions.

Figure 15: (a) Mean cell throughput under different soft reuse power ratios. (b) The relative contributions to the mean from the ASB and BSB. (HEXAGONAL Data, VF MCS Codeset)

Figure 15 shows the results using the VF MCS codeset on the hexagonal dataset. As can be seen, introducing MCS rates for SINRS greater than 18.5dB does not change the outcome of the experiment. On the contrary, the new curve accentuates the difference between reuse one and soft-reuse.

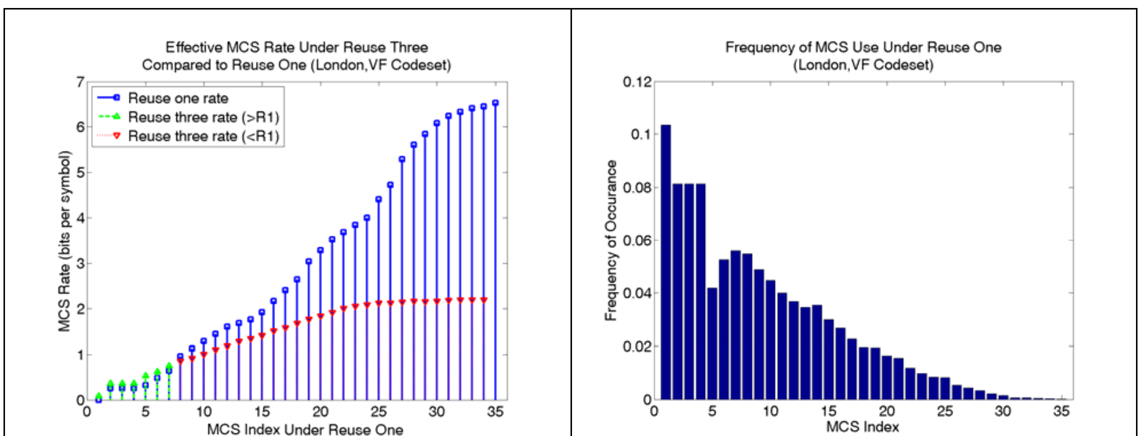

(a) Effective rate change when switching from reuse one to reuse three. **(b)** Relative proportional use of MCS schemes by UEs.

Figure 16: (a) Effective change in obtainable bitrate when switching from a reuse one interference scenario to a reuse three interference scenario. (b) Relative proportional use of MCS schemes across the cell (London, VF Codeset).

Figure 16 compares the MCS rate under reuse one with the effective MCS rate under reuse three, as before. As can be seen, the proportion of the cell in which reuse three compensates for the loss in bandwidth is much less than it was under the 3GPP MCS codeset.

Reuse three compensates for the loss in bandwidth in 47% of the cell area, with an average MCS rate improvement of 0.24 bits per symbol. Over the remaining 53% of the cell area, reuse three does not compensate for the loss in bandwidth, and the rate is reduced on average by 0.78 bits per symbol. The overall effect is that reuse one dominates. The effect is repeated in the non-hexagonal case.

The reason that this behavior is observed, is due to the slope of the MCS lookup curve. Switching to reuse three from reuse one causes a right shift in the SINR distribution, as shown in Figure 14.

This corresponds to a left shift in MCS lookup curve relative to reuse one. This is shown for the VF codeset in Figure 17**a**.

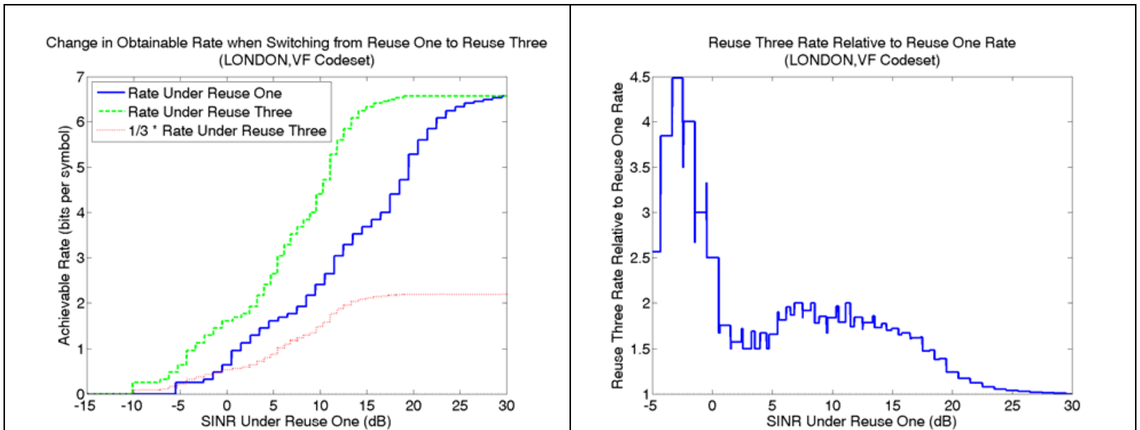

(a) Obtainable bitrate change	**(b)** Obtainable bitrate improvement

Figure 17: (a) Change in obtainable bitrate between reuse three and one. (b) Improvement in obtainable bitrate.

Figure 17**a** shows how although the rate is improved by a large extent by switching to reuse three, when compensated for the bandwidth reduction, a gain is only seen over a small region of SINR. Most of the gain seen can be attributed to giving throughput to UEs that formerly obtained none, and the rest is obtained by moving UEs away from the flat region seen at the base of the Reuse One curve. In general, a gain is likely to be seen, whenever the gradient of the rate curve under an SINR shift improves by a large enough extent.

Figure 17**b** plots the improvement in reuse three rate relative to reuse one rate. Anything greater than a factor 3 improvement in bitrate means that reuse one benefits.

It should be clear from this observation that the areas of the cell that benefit from coordinating interference are determined by three things:

1. The change in SINR distribution bought about by the coordination.

2. How this corresponds to a change in acheivable MCS rate.

3. The relative proportion of UEs which use each MCS rate.

2.2.4 Synopsis

Experiments were performed to assess the mean cell gain associated with applying soft reuse to simulated cellular networks based on hexagonal and realistic data. Results were presented for both a 3GPP reference SISO MCS lookup curve and a Vodafone obtained MIMO MCS lookup curve. In all cases the mean cell throughput was found to be better under reuse one when compared to any other soft reuse scheme tested including reuse three.

In conclusion: soft reuse provides no benefits over, and to the contrary performs worse than, reuse one, in the mean case.

2.3 Fixed Bitrate Scheduling

2.3.1 Motivation

In the last section, the conclusion was drawn that soft-reuse, as implemented here, provides no benefit in the mean case. These results, whilst illustrative, does not reflect the typical conditions under which a UE is served or under which a network operates. To address this, in this section, the same range of soft reuse schemes are examined for a fixed rate traffic service.

2.3.2 Procedure

It is assumed that a number of UEs desire data and that they all have the same fixed rate requirement. In the following plots this requirement will be changed, keeping the number of UEs constant, so that the contestation for resources changes. The fixed rate scheduler used operates as outlined in Procedure 2.

```
ForEach (UE in Scheduling Order) : Do
    Allocate VRBs from the BSB until either:
        No VRBs remain
        OR UE is satisfied.
    Allocate VRBs from the ASB until either:
        No VRBs remain
        OR UE is satisfied.
    Increment number of UEs satisfied if required.
Done
```

Procedure 2: Fixed rate scheduling algorithm

Observe that this scheduling procedure favours the UE that is visited first. To take a look at how results vary as the scheduling order is changed, three scheduling order strategies were tested here:

1. GREEDY - The UEs are scheduled according to wideband SINR in descending order from best to worst.
2. RANDOM - The UEs are scheduled in random order.
3. LEFTIST - The UEs are scheduled according to wideband SINR in ascending order from worst to best.

The number of UEs was fixed at 50. Results were averaged over 1000 random user drops, and over all shadowing seeds or cell positions as before. The target rate was varied as specified in Table 4.

Bitrate Index	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Bitrate (kbps)	150	100	90	80	70	60	50	40	35	30	25	20	15	10	5

Table 4: Fixed rate scheduling bitrate targets (kbps).

2.3.3 Results

It would over complicate things to present results for all bitrates tested so in the following sections, three load cases are examined, followed by an overview to extract general trends as bitrate is varied. The three load cases are overload: where the number of UEs satisfied falls far short of the number scheduled; transition: where the number of UEs satisfied is verging on being the number of UEs scheduled, and underload, where all UEs that meet the minimum data rate SINR are satisfied.

2.3.3.1 Overloaded cell

For the overload case, consider the target rate of 150kbps. At this rate it is not possible to support all 50 UEs. Figure 18 shows the results.

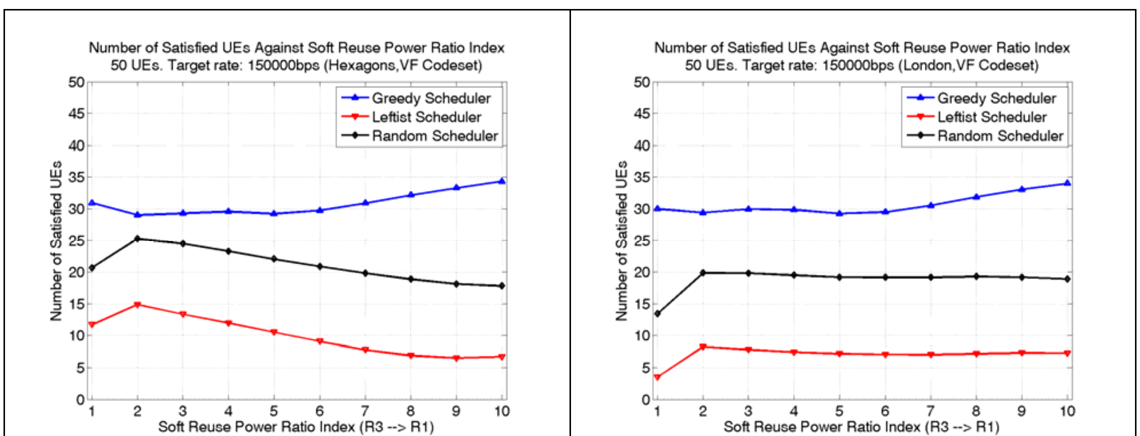

Hexagonal dataset	London dataset

Figure 18: Fixed rate scheduling capacity for an overloaded cell where each UE has a target bitrate of 150kbps.

For both the Hexagonal and London datasets, when the cell is overloaded, the greedy scheduling strategy significantly outperforms the other two scheduling strategies for all soft reuse power ratio points, and the random scheduler outperforms the leftist scheduler for all soft reuse power ratio points ($p < 0.05$ in all cases). For the Hexagonal dataset, and using the greedy scheduler, the number of

UEs satisfied under reuse one is about 13% more than the number satisfied under reuse three ($p < 10^{-6}$). For the London data set, reuse one satisfies about 11% more UEs than reuse three ($p < 0.0001$).

The results imply that in overload situations, and for fixed rate services, reuse one outperforms any soft reuse scheme including reuse three.

The reason that the leftist and random scheduling schemes satisfy more UEs as the reuse scheme tends towards reuse three for the Hexagonal dataset, is because the BSB power is close to that of reuse three, yet the additional bandwidth on the ASB, albeit poor quality, allows for a few extra UEs to be supported.

It is unlikely that any practically deployed network will operate at the level of contention for fixed services suggested in Figure 18, yet since the results are independent of bandwidth, it is feasible to conceive of business models where the situation might arise for a part of the band. That aside, the next important step is to examine lower target rates, where the number of supported UEs approaches or exceeds the maximum possible.

2.3.3.2 Transitioning cell

As the target bit rate is reduced, the number of UEs satisfied approaches the number of UEs served, that is, the cell transitions from not being able to serve all UEs, to being able to serve all UEs. Figure 19 shows the results, on both the hexagonal and london datasets for target bit rates selected close to the transition point in each case.

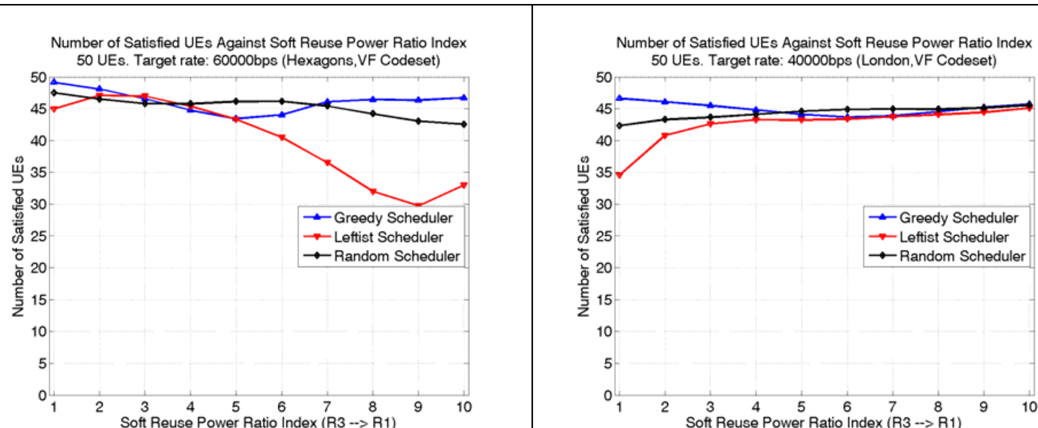

Hexagonal dataset (VF Codeset)	London dataset (VF Codeset)

Figure 19: Fixed rate scheduling capacity for target bitrates at the point before peak UE satisfaction.

In each case about 90% of UEs are satisfied. The general trend observed, is that as the transition point is approached, scheduling strategy differences are reduced.

However, in the hexagonal case, under the greedy scheduler, it now happens that reuse three serves about 5% more UEs than reuse one ($p < 10^{-6}$). In the London case also, reuse three shows a non-significant improvement over reuse one of about 2% ($p = 0.18$).

To make two final observations, note that the London dataset lags behind the Hexagonal dataset in that less UEs can be supported for a given target bitrate. This is because, with reference to Figure 14, the London dataset CDF is left shifted and has a longer tail in comparison to the Hexagonal dataset. Also note that the low SINR UEs suffer less improvement when switching to reuse three, this supports subjective statements such as ``users at the cell edge ... suffer the most from inter-cell interference'' as stated in [26].

2.3.3.3 Underloaded cell

Figure 20 characterizes the underloaded cell case, where all UEs can be satisfied.

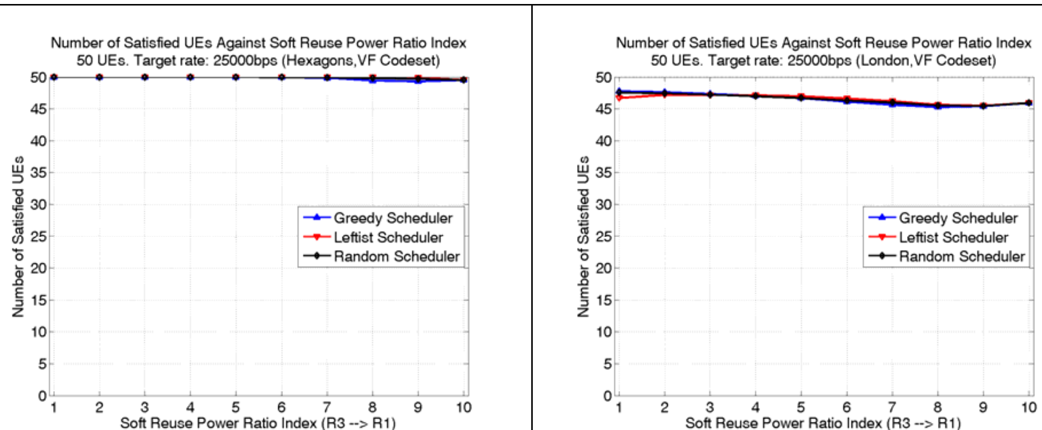

Hexagonal dataset (VF Codeset)	**London dataset (VF Codeset)**

Figure 20: Fixed rate scheduling capacity for an under-loaded cell where each UE has a target rate of 25kbps.

Note that for both the hexagonal and London datasets, and especially so in the latter case, the cell on average can never satisfy all UEs because, again with reference to Figure 14, a small percentage of UEs have SINRs less than the minimum required of $-6.5dB$.

With this caveat in mind, observe that for the London dataset reuse three satisfies about 4% more UEs than does reuse one ($p < 0.02$). And for the hexagonal dataset, reuse three satisfies about 1% more UEs than reuse one ($p < 10^{-10}$). This is easy to explain by observing the right shift that reuse three imparts to the SINR CDF as demonstrated in Figure 14.

2.3.3.4 Bitrate dependency overview

The results so far suggest that for overload situations reuse one performs upto 10% better than reuse three, and for underload situations, where all UEs can be satisfied, reuse three performs upto 5% better than reuse one.

These results indicate that the best soft reuse power ratio will be dependent on the loading conditions.

Figure 21 and Figure 22 plot, for each target bitrate, the maximum number of UEs supported under the optimal soft reuse power ratio for that bitrate. They also plot, as an index relative to Table 3 what that optimal soft-reuse power ratio is. Figure 21 plots this information for the hexagonal dataset and Figure 22 plots this information for the London dataset.

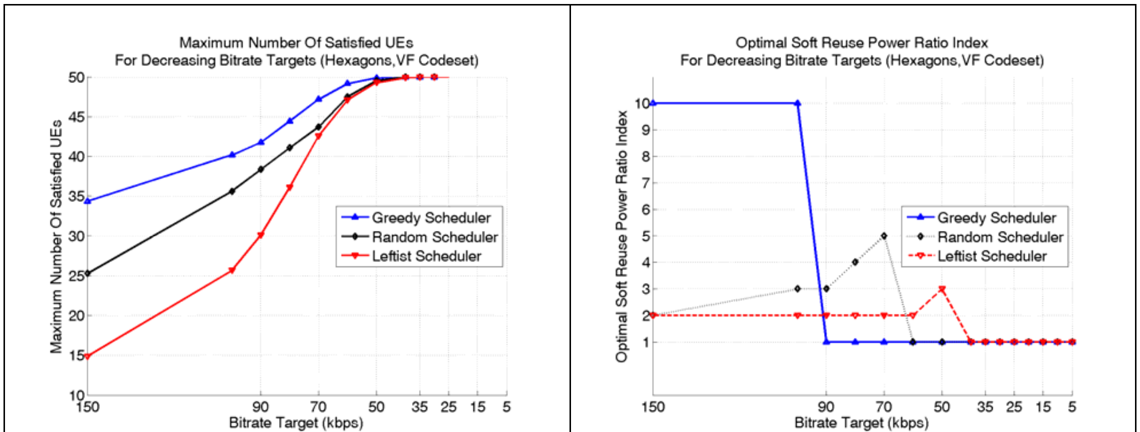

Maximum number of supported UEs (Hexagons, VF Codeset)	Power ratio index that achieves maximum in plot to the left (Hexagons, VF Codeset)

Figure 21: Optimal soft-reuse power ratio and maximum number of satisfied UEs for each target bitrate. The optimal soft-reuse power ratio is a number between 1 and 10, which corresponds to the soft-reuse power ratios studied, as shown in Table 3. This plot is for the Hexagonal dataset and the VF Codeset.

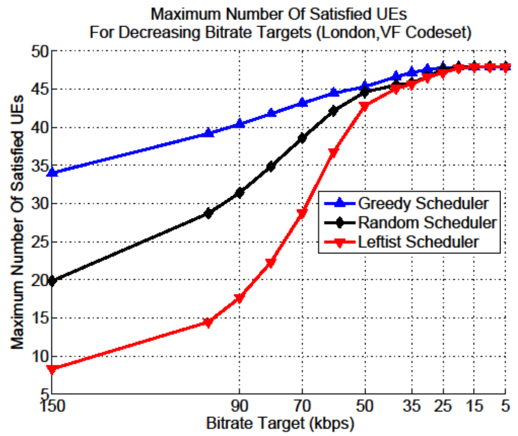	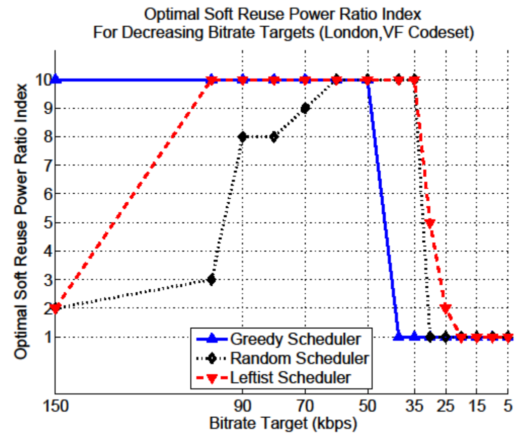
Maximum number of supported UEs (London, VF Codeset)	**Power ratio index that achieves best performance in plot to the left (London, VF Codeset)**

Figure 22: Optimal soft-reuse power ratio and maximum number of satisfied UEs for each target bitrate. The optimal soft-reuse power ratio is a number between 1 and 10, which corresponds to the soft-reuse power ratios studied, as shown in Table 3. This plot is for the London dataset and the VF Codeset.

Consider first, the maximum number of UEs that can be satisfied for each target bitrate. As the target bitrate is reduced, the number of satisfied UEs increases. This is true for both the hexagonal and London datasets and is expected.

Furthermore, the greedy approach to scheduling always outperforms or differs insignificantly from, the other two approaches.

Now consider the right-hand part of each plot, where the optimal soft-reuse power ratio is plotted. This is the soft-reuse power ratio under which the number of satisfied UEs is maximized, for the given target bitrate.

Considering only the greedy scheduling strategy, since the other strategies perform worse, it should be clear that in both cases, the optimal power ratio changes from 10 (reuse three) to 1 (full reuse) as the bitrate target is reduced.

For none of the bitrate targets, do any of the intermediate soft-reuse power ratios perform best.

Given then, that either reuse one performs best, or reuse three does. Figure-17 shows how reuse three performs relative to reuse one (as a fraction).

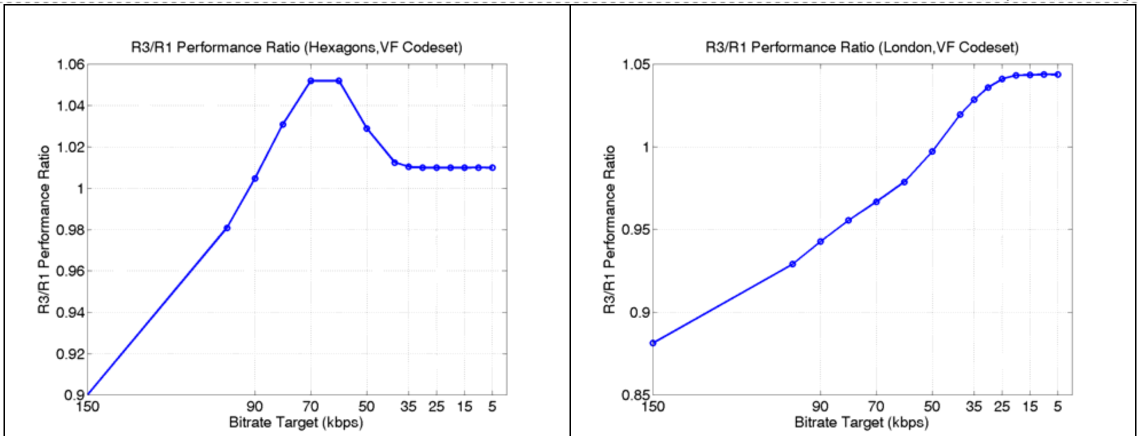

Hexagonal dataset (VF Codeset)	**London dataset (VF Codeset)**

Figure 23: The performance of reuse three relative to reuse one.

In these plots, when the relative performance measure for reuse three is less than one, this means that reuse three performs worse than reuse one. When it is greater than one, this means that reuse three outperforms reuse one.

The point at which each plot crosses the threshold of one is of course equal to the point at which the optimal scheme in Figure 22 changes from reuse one to reuse three.

Considering the hexagonal dataset. When reuse one performs better it performs a maximum of $\approx 10\%$ better. When reuse three performs better, it performs a maximum of $\approx 5\%$ better, and this maximum occurs close to the switch point. For lower target bitrates, this gain reduces and seems to stabilize at $\approx 1\%$.

Taking the London dataset, reuse one performs better, it performs a maximum of $\approx 13\%$ better. It performs better than reuse three for more bitrate targets than in the hexagonal case. When reuse three performs better, it performs a maximum of $\approx 4\%$ better. This gain seems to stabilize around this value for lower bitrate targets.

2.3.4 Synopsis

- The greedy scheduling strategy performs better than either the random or the greedy scheduling strategies. It always pays to start with the UEs with the best SINR, and give the best resource to them.
- For high target bitrates, where not all UEs can be satisfied, reuse one performs upto 10% better than any other soft-reuse power ratio, including reuse three.
- For low target bitrates, when all UEs can be supported, reuse three performs upto 4% better than any other soft-reuse power ratio, including reuse one.

The results suggest that small gains in the maximum number of satisfiable UEs can be acheived by using reuse three, for low bitrate targets.

The reason for this is because under reuse three, the SINR curve is right shifted, so it less often occurs that a UE cannot be satisfied at all, and it less often occurs that very low bitrates are needed.

None of the intermediate soft-reuse power ratio performed better than either reuse one or reuse three. It seems reaonable to hypothesize that close to the transition area, where the optimal scheme changed from reuse one to reuse three, that they might however.

If this transition area were to be investigated at a higher resolution, perhaps then some of the intermediate soft reuse power ratios would be the optimal choice, but even if this were the case, it is likely the gains would not exceed the gains of reuse three when it performed best.

Given that the simulations in this section are static and monte carlo, it is difficult to quantify whether the small gains seen here would occur in a dynamic simulation. It seems more likely that the performance would be dominated by other factors, and that the small gains observed here, which might be around 2-3% in practice, are not worth the coordination effort.

Furthermore, relying on these gains depends on the greedy scheduling strategy, and thus SINR prediction. If this were to fall away, it is unlikley the gains would be seen at all.

In conclusion, for fixed rate services, under limited conditions, small gains can be observed in employing reuse three over reuse one. But these gains are not really large enough to warrant the coordination effort.

2.4 Throughput Fair Scheduling

2.4.1 Motivation

2.2 demonstrated that reuse three bought about net gains only for limited regions of the cell. Specifically, only those UEs having low SINRs in the reuse one case, were shown to obtain a net benefit when switching to reuse three. This was due to the shape of the MCS curve. At the lower end of the MCS curve, switching to reuse three bought about a change in rate that was greater than a factor of three, but at the upper end of the MCS curve it did not. This was because the lower part of the MCS curve was the only part with a gradient greater than three. Under the relative porportional use of MCS schemes observed, this produced the mean result that reuse one outperformed reuse three.

These findings are now examined for multi-user allocations via the inspection of allocations which are biased toward different parts of the MCS lookup curve. In order to make sense of such allocations, they are related to a more understandable metric: Jain's fairness index [22] as applied to UE throughputs.

The Jain's fairness index function is defined in Equation-2.1 below:

$$Jains(x) = \begin{cases} 1 & x = 0 \\ \left[\sum_{i=1}^{N} x_i\right]^2 \cdot \left[n \cdot \sum_{i=1}^{N} x_i^2\right]^{-1} & otherwise \end{cases} \tag{2.1}$$

Where x_i is the quantity of resource associated with UE i. The vector $x = [x_1\ x_2\ ...\ x_N]$ can refer to any measurable resource. When the resource is throughput, the fairness is called here throughput fairness, and when the resource is allocated VRBs, the fairness is called here resource fairness. The domain and codomain of the Jain's fairness index function are shown below:

$$Jains : \mathrm{R}^N \rightarrow [\frac{1}{N}, 1]$$

The useful properties of this fairness index are that: (i) An equal distribution of resource results in a fairness of 1 (ii) Giving all the resource to one UE results in a fairness of 1/N (iii) Giving M UEs an equal share of the resource results in a fairness of M/N.

These intuitive properties of Jain's fairness index make it appealing to use but it suffers from one critical drawback: in the context of cellular networks, total and percentile throughputs are heavily influenced by exactly which UEs are given throughput. Yet Jain's fairness will evaluate to 1/N irrespective of which of the N UEs is given all the resource. More generally, Jain's fairness index function is non-injective and thus suffers from the problem that it is difficult to interpret the meaning of any given fairness result.

For this reason, the allocations produced by the algorithm to follow are very carefully controlled so that evaluated throughput fairness makes intuitive sense given the context.

2.4.2 Procedure

A sweep of throughput fairness values were examined between the following three successive waypoints:

1. Throughput fairness $=1/N$ (Lowest bitrate biased).

   The UE with lowest bitrate is given all the resource and thus all the throughput. This produces an allocation heavily biased toward the lower end of the MCS lookup curve, and should thus favor the use of orthogonal schemes over reuse one.

2. Throughput fairness $=1$.

   All UEs receive the same throughput. It is unknown apriori how this will influence use of MCS schemes, but since low bitrate UEs will be given proportionately more resource than high bitrate UEs, this may bias the mean result in favor of orthogonal schemes over reuse one.

3. Throughput fairness $=1/N$ (Highest bitrate biased).

   The UE with highest bitrate is given all the resource and thus all the throughput. This produces an allocation heavily biased toward the upper end of the MCS lookup curve, and should thus favour reuse one over orthogonal schemes.

The fairness sweep starts with a throughput fairness of $1/N$, rises to a throughput fairness of 1, and then falls to a throughput fairness of $1/N$. And although the same throughput fairness values are visited either side of the peak, these values have carefully controlled meanings: the (un)fairness is either biased toward the UE with lowest bitrate or biased toward the UE with highest bitrate.

In order to produce this sweep, two items are required: firstly, the allocation vector for the throughput fair $=1$ point described above must be determined; and secondly, a means of interpolating between this and the end waypoints, so that intermediate throughput fairness values can be examined must be obtained.

To address the first item, let $\hat{a}=[a_1,a_2,...,a_N]$ denote the allocation vector which results in a throughput fairness of 1, where a_j is the fraction of resource allocated to UE j and $\sum_{j=1}^{N} a_j = 1$. Such that, given N UEs with obtainable bitrates $b=[b_1,b_2,...,b_N]$, it follows that $Jains(\hat{a}.*b)=1$, where .* denotes elementwise multiplication. To construct $\hat{a}$ UE j should receive the following fraction a_j of the resources:

$$a_j = \left[b_j \cdot \sum_{i=1}^{N} b_i^{-1} \right]^{-1}$$

To address the second item, a means of modifying $\hat{a}$ is required, which is capable of biasing the allocation toward either of the end waypoints (where either the lowest bitrate or highest bitrate UE receives all resources). With this goal in mind, assume that a modifier vector m_T^{β} exists, such that $|\,Jains(\hat{a}.*m_T^{\beta}.*b)-T\,|<\varepsilon$ for any target throughput fairness T, target accuracy ε, and target bias direction $\beta = \{-,+\}$, where $-$ signifies that the allocation is biased toward the case where all resources are given to the UE with worst bitrate, and $+$ signifies that the allocation is biased toward the case where all resources are given to the UE with best bitrate. The aim is to find m_T^{β} for all target T and for both bias directions.

Notice that a brute force search for a given m_T^{β} would not be possible, since it is very difficult to establish apposteriori whether a given allocation displays the correct bias, since so many different random allocations will map to the same fairness value. To this end m_T^{β} are obtained by searching through a set of candidate modifiers that are known to display the desired bias. For this purpose, the exponential function has properties which make it ideal. To see this, consider the discrete exponential function $e^{\alpha x}$ for the vector $x = [-1.0\ -0.5\ \ 0.0\ \ 0.5\ \ 1.0]$. This function is plotted in Figure 24 for three different values of α.

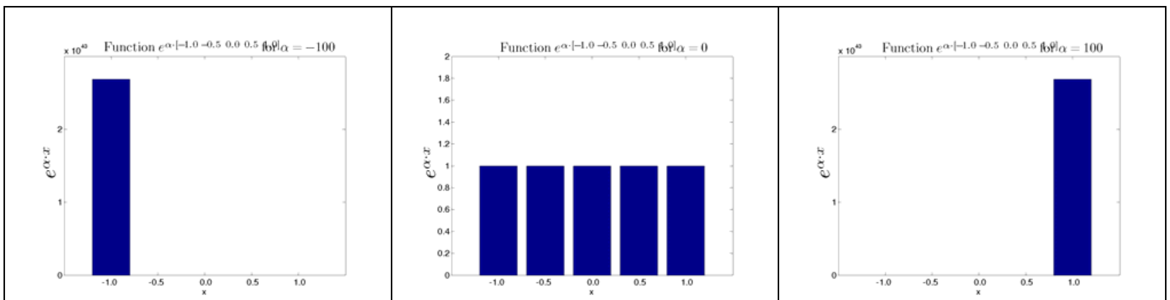

$\alpha = -100$	$\alpha = 0$	$\alpha = 100$

Figure 24: The function $e^{\alpha x}$ for $x = [-1.0\ -0.5\ \ 0.0\ \ 0.5\ \ 1.0]$ and $\alpha = -100$, 0, and 100.

Imagine that these three functions are applied as candidate modifiers to $\hat{a}$ via elementwise multiplication, and the resultant vector normalized so that it represents a valid allocation for 5 UEs. It should be clear that the three α values displayed above, will produce the vectors $m_{1/N}^{-}$, $\hat{a}$, and $m_{1/N}^{+}$ whenever the 5 UEs in b are sorted into ascending order of bitrate. Empirically, all m_T^{β} for all desired targets could be found given the appropriate α input.

- For a bias direction $\beta = -$, and target fairness T, m_T^{-} can be found by performing a sweep of α values less than zero.
- For a target value of $T = 1$, $\alpha = 0$ produces the correct modifier vector.
- For a bias direction $\beta = +$, and target fairness T, m_T^{+} can be found by performing a sweep of α values greater than zero.

Of course the function $e^{\alpha x}$ needs in practice to have as many discrete steps as there are UEs, and so in the experiments that follow x was set to $[-1 - \frac{11}{12} - \frac{10}{12} \ldots 0\ \frac{1}{12}\ \frac{2}{12} \ldots 1]$ to give 25 discrete steps.

Using this search procedure, for any given bitrate vector b, allocations could be found for the throughput fairness targets shown in Table 5 and $\varepsilon = 5\times10^{-4}$:

Worst UE Bias	$1/N^-$	0.1^-	0.2^-	0.3^-	0.4^-	0.5^-	0.6^-	0.7^-	0.8^-	0.9^-
No Bias	1									
Best UE Bias	0.9^+	0.8^+	0.7^+	0.6^+	0.5^+	0.4^+	0.3^+	0.2^+	0.1^+	$1/N^+$

Table 5: Throughput fairness targets examined. Throughput fairness targets with a negative superscript are biased toward giving all resources to the UE with worst bitrate. Throughput fairness targets with a positive superscript are biased toward giving all resources to the UE with best bitrate.

To illustrate the application of this procedure, Figure 25 plots the obtained throughput fairness, resource fairness, and throughput for each target fairness, for a random vector of 25 bitrates drawn from $(0,25]$.

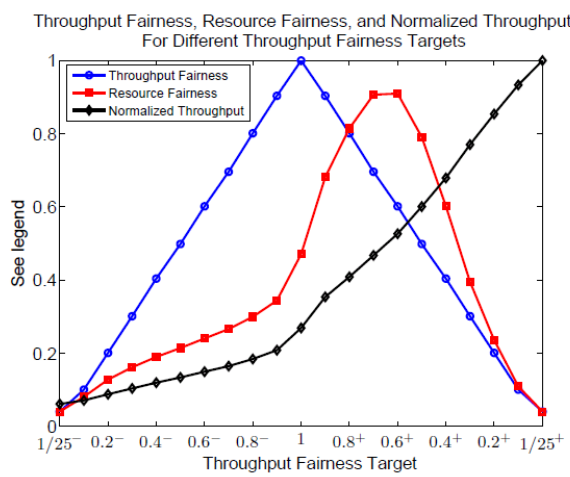

Figure 25: The evolution of total throughput, throughput fairness, and resource fairness, for a sweep of throughput fair targets for 25 randomly selected UEs.

The figure demonstrates how the procedure is able to match throughput fairness targets. For low fairness values biased toward the UE with worst bitrate, throughput is very low. As the throughput fairness is increased toward a value of 1, throughputs increase and resource fairness increases. Beyond a fairness target of 1, the fairness is reduced toward the direction of the UE with best bitrate, and thus the throughput increases still further. Resource fairness then peaks and beyond this, it must reduce in order to give the best bitrate UE proportionately more resource, and thus decreases to $1/N$ as the throughput comes to its peak.

Given this procedure for obtaining an allocation vector matched to a throughput fairness target, Procedure 3 outlined the procedure used to obtain the results plotted in the next section.

```
ForEach (Soft Reuse Power Ratio) : Do
    ForEach (Shadowing Seed) : Do
        For (UE Drop Count) In (1 to 1000) : Do
            Select 25 UEs at random from service area
            ForEach (Throughput Fairness Target) : Do
                Find an appropriate allocation vector
                A <- Total throughput
                B <- 12th percentile throughput
                C <- 8th percentile throughput
                D <- 4th percentile throughput
                E <- Throughput fairness
                F <- Resource fairness
            Done
        Done
        Compute mean of A,B,C,D,E,F over all UE Drop Counts
    End
    Compute mean of A,B,C,D,E,F over all Shadowing Seeds
End
```

Procedure 3: Process for obtaining throughput fair experiment results

2.4.3 Results

Figure-20 plots the results for the two extremes of unfairness; the throughput fairness target $1/25^-$ where the worst bitrate UE gets all resource, and the throughput fairness target $1/25^+$ where the best bitrate UE gets all the resource.

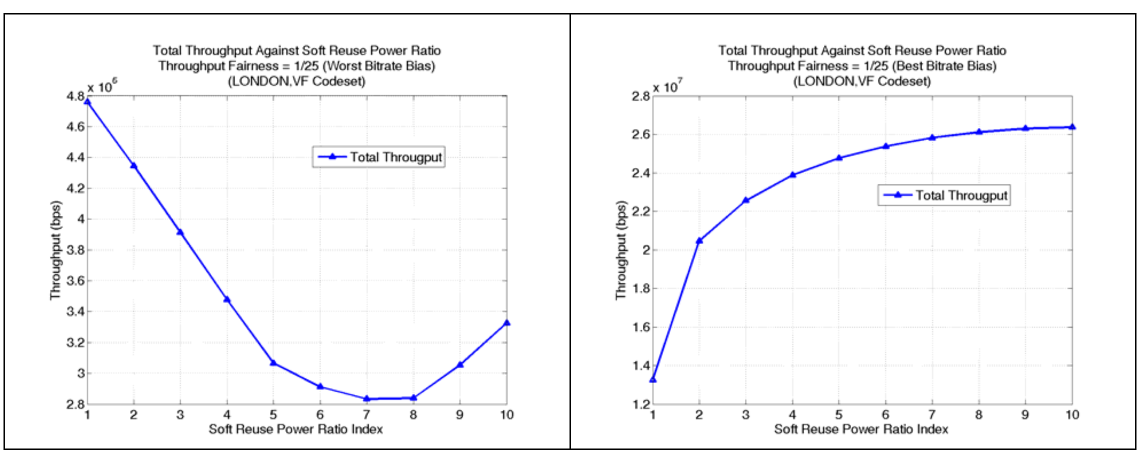

Fairness Target $1/25^{-}$, Worst Bitrate Bias	Fairness Target $1/25^{+}$, Best Bitrate Bias

Figure 26: Performance under different soft reuse power ratios for two resource allocations which both evaluate to a throughput fairness of $1/25$. The difference in the allocations is that in the worst bitrate bias case, all the resources are given to the UE with worst bitrate, whereas in the best bitrate bias case, all the resources are given to the UE with best bitrate.

The left plot corresponds to the case where all resources are given to the UE with the worst bitrate. In this case, the best performing soft-reuse power ratio, was reuse three, which outperformed reuse one by $\approx 43\% (p = 0.057)$. The right plot corresponds to the case where all resources are given to the UE with best bitrate. In this case the best performing soft-reuse power ratio is reuse one, which outperforms reuse three by $\approx 99\% (p < 1\times10^{-13})$.

This demonstrates clearly that scheduling strategy and thus resource allocation, can have a critical effect on the relative performance of different soft-reuse power ratios, even when the throughput fairness of those allocations evaluate to the same number.

The result shown is easy to explain. It was demonstrated in Figure-9 that reuse three produces more than three times the throughput of reuse one only for points which obtain a low bitrate under reuse one. In the left plot, the bias is towards this region, in the right plot it is away from this region.

Figure 27 shows the total throughput and percentile throughputs, across all soft reuse power ratios examined, for a throughput fairness target of 1.

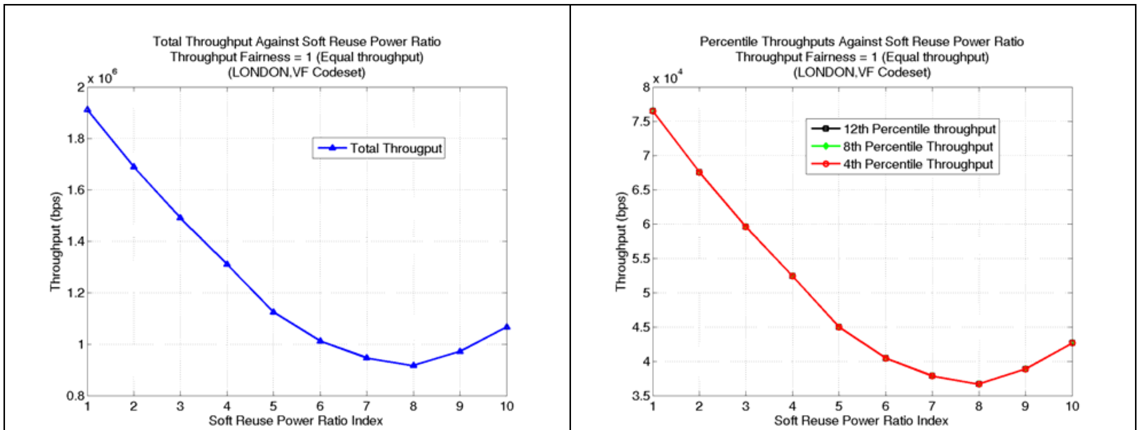

Total throughput	Percentile throughputs

Figure 27: Total and percentile throughputs under different soft reuse power ratios for a throughput fairness target of 1.

Considering the total throughput, reuse three gives $\approx 80\%$ greater throughput than reuse one $(p = 0.0026)$. Again this can be understood since in the throughput fair case, proportionately more resource must be given to the low bitrate UEs, which means that in Figure 27 more UEs are in the region where reuse three compensates for the loss in bandwidth.

The percentile throughputs shown are simply $1/25$ of the total througput, and are equal, since every UE gets the same throughput.

The results shown so far demonstrate that at a throughput fairness of $1/25^-$ (biased toward worst bitrate UEs) reuse three performs best, that at a throughput fairness of 1 reuse three again performs best, but at a throughput fairness of $1/25^+$ (biased toward best bitrate UEs), reuse one performs best. Clearly then, at some point between fairness values of 1 and $1/25^+$, the best performing soft reuse power ratio scheme must switch from reuse three to reuse one.

The point at which this occured in these experiments, given the discretization of fairness targets, was when the target fairness was 0.3^+ (biased toward best bitrate UE). Figure 28 shows the total and percentile throughputs for this point.

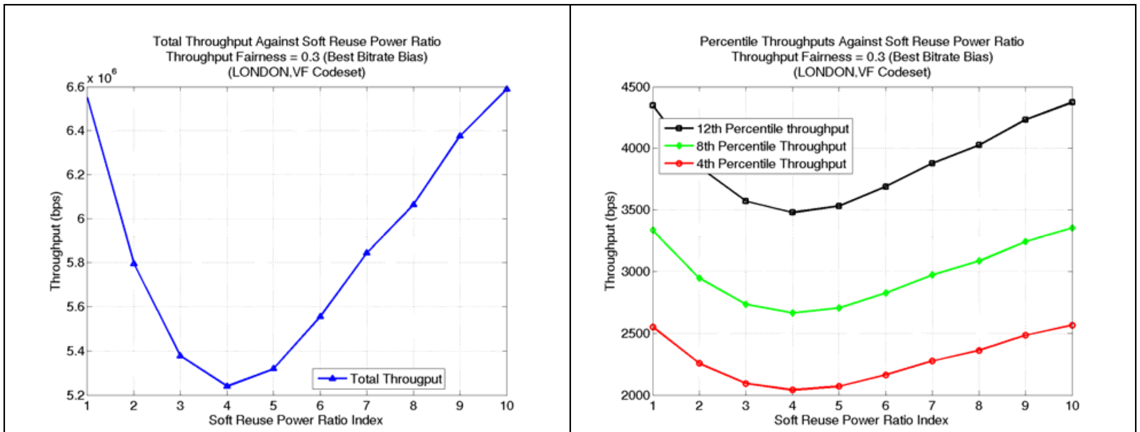

Total throughput	Percentile throughputs

Figure 28: Total and percentile throughputs for a Jain's throughputs fairness of 0.3 (high SINR bias).

An observation worth noting is that the curve is convex, this makes it difficult for any scheme other than reuse one or three to have optimal performance. To examine this assertion, Figure 29 plots the best obtainable, among all soft reuse power ratios, total and percentile throughputs for each throughput fairness target.

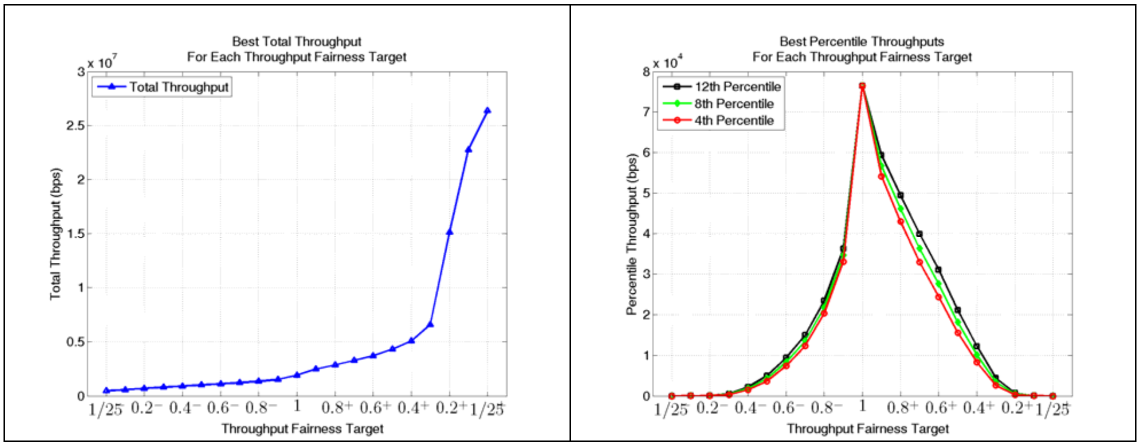

Total throughput | Percentile throughputs

Figure 29: Best total and percentile throughput for each Jain's throughput target.

The figure shows how total throughput goes from a minimum when all throughput is given to the worst bitrate UE, and increases as the sweep through throughput fairness targets is undertaken to reach a maximum when all the throughput is given to the best bitrate UE. The percentile throughputs, on the the other hand, peak at the throughput fair point and decay either side of it.

Figure 30 **(a)** plots the indices of the soft-reuse power ratios responsible for producing the throughputs shown in Figure 29.

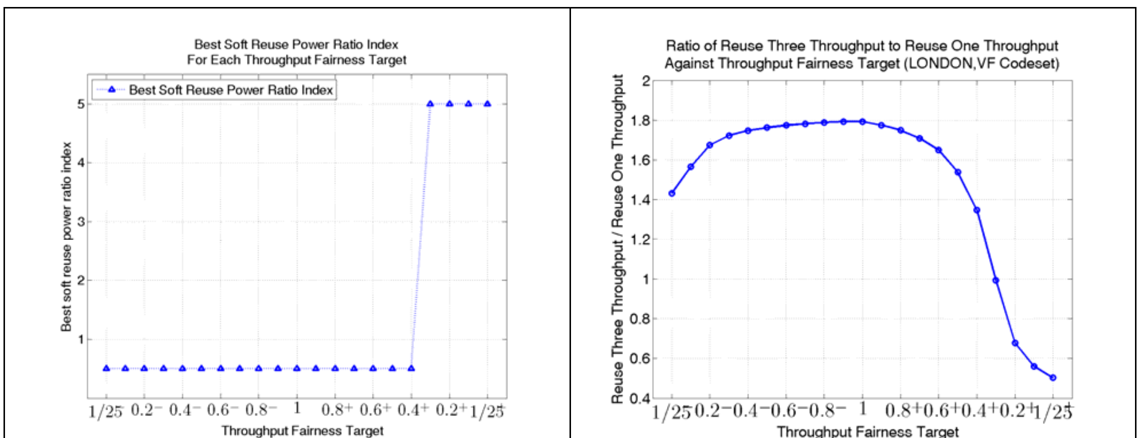

(a) Index of Best Soft Reuse Power Ratio | **(b)** R3 Throughput / R1 Throughput

Figure 30: (a) **Best soft reuse power ratio index for each throughput fairness target, for total and percentile throughput measures.** (b) **Ratio of reuse three througput to reuse one througput.**

Figure 30 **(a)** shows how the best soft reuse power ratio to use switches immediately from reuse three to reuse one, as soon as the bias toward low bitrate UEs is sufficiently lifted. It demonstrates that no intermediate soft reuse scheme ever performed the best.

Figure 30 **(b)** plots the ratio of reuse three to reuse one throughput for each throughput fairness target. For a large range of throughput fairness target, the reuse three total throughput exceeds that of reuse one by more than 50%.

2.4.4 Synopsis

The results shown in this section support the predictions made in Section 2.2 regarding the proportional use of MCS schemes, these can be summarized as follows:

1. When the throughput distribution is biased toward low bitrate UEs, the proportional use of MCS schemes is shifted away from the mean case such that reuse three outperforms reuse one by upto 80%.

2. When the throughput distribution is biased toward high bitrate UEs, the porportional use of MCS schemes favours reuse one, which outperforms reuse three by upto 50%.

Thus the results illustrate that multiuser allocations can be biased in a direction that favors either reuse one or reuse three. This is directly related to the interplay between MCS lookup curve and the SINRs distribution under the compared reuse schemes.

It is worth noting that no intermediate soft reuse power ratio outperformed reuse three, which meant that enforcing perfect frequency orthogonality was the best solution in the cases where reuse one did not perform best.

2.5 Discussion

When any interference coordination scheme is applied, it should result in a right shift of at least part of the SINR CDF, as was demonstrated for reuse three in Figure-8. If a right shift is not observed for any part of the CDF, then no improvement in throughput will be obtained. Assuming that bitrate is monotonically increasing with SINR.

The improvement in bitrate obtained, integrated over the entire SINR CDF, must exceed any losses incurred over the CDF by the introduction of the interference coordination scheme as illustrated in Figure 14.

The outcome is not only determined by SINR CDF and MCS bitrate curve, but also the relative proportional use of each MCS scheme. This can be changed by modifying the scheduler. Scheduler policy usually follows operator policy regarding priorities that should be assigned to different UEs. And so it is important that when interference coordination schemes are examined, that they are examined under the regime of the operator policy that is of interest. Otherwise the results may be different in practice than in theory.

In the Monte Carlo simulations examined here, only the fixed rate scheduling approach is comparable to something like the services seen in real networks. Under this traffic regime, the results showed at most a 4% advantage by statically applying reuse three in a realistic network setting. This advantage is smaller than the static defecit introduced into the cell, ie the proportion that cannot on average be served.

Given this result, and the possibility that other traffic regimes may need to be supported for which reuse three is not favored, it would be dangerous to recommend that networks be deployed in reuse three. This is convenient, since the coordination cost of reuse three is also under scrutiny.

It would make more sense to coordinate the resource usage of individual UEs on a limited scale, and only in those cases where a gain is likely to occur, given assumptions about what will happen to the bitrates of those affected UEs before and after coordination.

For example, one could coordinate the scheduling of pairs of UEs that are both using an MCS scheme in the critical region and make their resources orthogonal.

Monte carlo simulations are limited and only express what happens on average over the area of the cell. The biggest problem is that they neglect the time and frequency dependent gains that be obtained through differential scheduling strategies.

Note that the results have been presented for reuse three as a comparison to reuse one but would be applicable to any interference coordination mechanism that sacrifices resources in favour of an SINR CDF right shift.

2.6 Conclusion

Scheduling strategy, SINR CDF, and MCS lookup curved employed are the most important factors in determining whether a given ICO scheme will manifest a benefit.

Static soft reuse schemes are not a good candidate for interference coordination in cellular networks because they behave differently depending on the traffic served and the user distributions. For a generic set of cells it would therefore make no sense to deploy static soft reuse unless it was absolutely clear apriori what kind of traffic, operator policy, SINR distributions, and user distributions are to be expected.

Given this, and the observation that interference coordination can provide gains under certain limited circumstances, it makes more sense to pursue ideas that pertain to identifying these particular circumstances and performing dynamic coordination of only those UEs who need it.

3 REFERENCES

[1]Digital Mobile Radio Towards Future Generation Systems. Cost 231 Final Report. Chapter 4: Propagation Prediction Models. http://www.lx.it.pt/cost231/docs/PDFfiles/Chap4_pdf.ZIP.

[2]S.M. Alamouti. A simple transmit diversity technique for wireless communications. *Selected Areas in Communications, IEEE Journal on*, 16(8):1451 –1458, Oct 1998.

[3]Alexander L. Stolyar and Harish Viswanathan. Self-organizing Dynamic Fractional Frequency Reuse for Best-Effort Traffic Through Distributed Inter-cell Coordination. In *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, pages 691–699, April 2008.

[4]Atoll 2.8 Brochure: RF Planning & Optimisation Software. http://www.forsk.com/-automne_modules_files/standard/public/p11_79991b7e0d5af29cefa7dc2e15026d03Atoll_2.8_HD.pdf.

[5]Douglas G. Altman. *Practical Statistics for Medical Research*. Chapman and Hall/CRC, 1990.

[6]Gabor Fodor and Chrysostomos Koutsimanis and András Rácz and Norbert Reider and Arne Simonsson and Walter Muller. Intercell Interference Coordination in OFDMA Networks and in the 3GPP Long Term Evolution System. *Journal of Communications*, 4(7):445–453, August 2009.

[7]Gary Boudreau and John Panicker and Ning Guo and Rui Chang and Neng Wang and Sophie Vrzic. Interference Coordination and Cancellation for 4G Networks. *IEEE Communications Magazine*, 47(4), April 2009.

[8]Huawei. 3GPP TSG RAN WG1 Meeting #41, R1-050507: Soft Frequency Reuse Scheme for UTRAN LTE, 2005.

[9]National institute of standards and technology, ieee 1588 website. accessed 17.12.09. http://-ieee1588.nist.gov.

[10]Texas Instruments. R1-051059: Inter-Cell Interference Mitigation for EUTRA, October 2005.

[11]Texas Instruments. R1-060067: Performance of Inter-Cell Interference Mitigation with Semi-Static Frequency Planning for EUTRA Downlink, 2006.

[12]John E. Rowe and M.D. Vose and Alden H. Wright. Reinterpreting No Free Lunch. *Evolutionary Computation*, 17(1):117–129, November 2009.

[13]3GPP TS 36.211 - V8.7.0 - Evolved Universal Terrestrial Radio Access (E-UTRA); Physical Channels and Modulation, May 2009.

[14]3GPP TS 36.942 - V8.2.0 - Evolved Universal Terrestrial Radio Access (E-UTRA); Radio Frequency (RF) system scenarios, May 2009.

[15]3GPP TS 36.214 - V8.7.0 - Evolved Universal Terrestrial Radio Access (E-UTRA) Physical layer – Measurements (Release 8), September 2009.

[16]3GPP TS 36.104 - V9.1.0 - Evolved Universal Terrestrial Radio Access (E-UTRA) Base Station (BS) radio transmission and reception (Release 9), September 2009.

[17]Markus Rupp and Christoph F. Mecklenbrauker. On Extended Alamouti Schemes for SpaceTime Coding. In *In Proceedings of 5th International Symposium on Wireless Personal Multimedia Communications (WPMC '2002), October 27-30, Hawaii, USA*, 2002.

[18]M.C. Necker. Towards Frequency Reuse 1 Cellular FDM/TDM Systems. In *Proceedings of the 9th ACM/IEEE International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM 2006)*, October 2006.

[19]R Bosisio and U Spagnolini. Interference Coordination Vs. Interference Randomization in Multicell 3GPP LTE System. In *Wireless Communications and Networking Conference, 2008. WCNC 2008. IEEE*, pages 824–829, 2008.

[20]R1-050738, Siemens, Interference mitigation Considerations and Results on Frequency Reuse, RAN WG1#42, London, UK, Aug./Sept. 2005.

[21]R1-051366, Aspects of Interference Mitigation by Coordination, Siemens, TSG-RAN WG1 Meeting #43, Seoul, Korea, November 2005.

[22]Raj Jain and Dah-Ming Chiu and W. Hawe. A Quantitative Measure Of Fairness And Discrimination For Resource Allocation In Shared Computer Systems. *DEC Labs Technical Report*, TR-301, 1998.

[23]Claude Elwood Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27:379–423, 623–656, 1948.

[24]Soumaya Hamouda and Sami Tabbane and Philippe Godlewski. Improved reuse partitioning and power control for downlink multi-cell OFDMA systems. In *ACM International Conference Proceeding Series; Vol. 196*, volume 196, Alghero, Italy, September 2006. ACM Press.

[25]Xueqin Xiang, Fuqiang Liu, and Yusheng Ji. Simulation based performance evaluation of ICI mitigation schemes for broadband wireless access networks. In *CNS '08: Proceedings of the 11th communications and networking simulation symposium*, pages 181–187, New York, NY, USA, 2008. ACM.

[26]Xiang Yikang, Luo Jijun, and Christian Hartmann. Inter-cell Interference Mitigation through Flexible Resource Reuse in OFDMA based Communication Networks. In *In proceedings of 13th European Wireless Conference (EW'07), Paris, France*, April 2007.

[27]F. Xiangning, C. Si, and Z. Xiaodong. An Inter-Cell Interference Coordination Technique Based on Users Ratio and Multi-Level Frequency Allocations. In *Wireless Communications, Networking and Mobile Computing (WiCom 2007)*, pages 799–802, September 2007.

APPENDIX A

A.1 Vodafone MCS Codeset Thresholds

Rate (bits per symbol)	SINR Threshold (dB)	Rate (bits per symbol)	SINR Threshold (dB)
0.25	-5.4	3.5236	13.5
0.3203	-2.4	3.6838	14.5
0.4805	-1.4	3.8439	15.5
0.6407	-0.4	4.0041	16.5
0.961	0.5	4.4045	17.5
1.1211	1.5	4.7248	18.5
1.2813	2.5	5.2854	19.5
1.4415	3.5	5.6057	20.5
1.6016	4.5	5.846	21.5
1.6817	5.5	6.0862	22.5
1.7618	6.5	6.2464	23.5
1.922	7.5	6.3265	24.5
2.1622	8.5	6.4066	25.5
2.4025	9.5	6.4466	26.5
2.6427	10.5	6.4866	27.5
3.0431	11.5	6.5267	28.5
3.2834	12.5	6.5667	29.5

Table 6: Vodafone MCS curve mapping between bitrate and SINR threshold

APPENDIX B

B.1 3GPP LTE Standards Restrictions

It is important to determine what, if any, restrictions the 3GPP LTE standard places on ICO solutions, since there may be some solutions which simply are not possible given the constraints. If this is true, and if such “illegal” solutions demonstrate excellent gains, then it would motivate the argument that the next standards release should incorporate the changes required to support those solutions. This section takes a look at the extant restrictions.

B.1.1 Sub-band transmit power

Many ICO solutions assume that transmit powers can be changed at some level of sub-band granularity, see for example [3].

The main problem with this assumption is that resource blocks in LTE do not only contain data symbols, but also contain pilot, control, and broadcast symbols. This raises the question: should the transmit power of these symbols also be changed? The answer is not straightforward.

Consider the pilot symbols. Pilot symbols in LTE are used to compute RSRP and to perform channel equalisation. Section 5.11 of [15] defines RSRP as the linear average of received power over all pilot symbols within the “considered measurement frequency bandwidth”. Unfortunately, this last “considered measurement frequency bandwidth” has been designated by the standards as an implementation decision, and not as a standards restriction. The only restriction imposed by the standards is that the UE must meet certain RSRP accuracy targets according to Section 6.5.4.1 of [16].

This means that UEs may be implemented which average over proper subsets of all pilot symbols. This introduces a problem, for if transmit powers have been changed over the measurement subset, outside of the UE’s knowledge, then its RSRP computations may no longer meet the accuracy requirements. If RSRP is inaccurate, despite being inviolate of the standards, then HO boundaries are liable to change and other problems could arise.

An obvious solution is to leave the pilot symbol TX power fixed at its normal value, however, since the positions in the frame of pilot symbols in adjacent cells are designed so as not to collide with each other, they inevitably collide with non-pilot symbols. Leaving the pilot symbol TX power at some fixed reference value could therefore result in the situation where low power data symbols are being interfered by the uncoordinated transmissions of high power pilot symbols.

So in light of this, it perhaps makes more sense to change the transmit powers of all symbol types. To overcome the UE RSRP problem, the standards could either mandate that UEs must average over all pilot symbols in the bandwidth, or the standards could define an efficient method of communicating power changes to UEs. The latter approach is less desirable than the first due to the additional overhead this signalling would incur. The former approach thus seems more reasonable, since then as long as the total TX power is kept constant, the average will always compute to the same value and therefore the RSRP will be unaffected.

Aside from the former problem, and assuming it can be overcome, it is interesting to ask the question: how much is transmit power allowed to vary across the signal bandwidth? To the best of our knowledge, it would appear that the standards do not explicitly prohibit the boosting of power within either certain regions of the bandwidth, individual VRBs, or indeed individual symbols. And there appears to be no fundamental technical reason, or inviolable law, against having a modest range of possible transmit powers that can be selected on a per symbol, per-VRB basis.

Even if the standards currently prohibited such power variation, the majority of frequency based interference coordination approaches published in the literature, assume it can be done. And therefore it is wise to assume here that it can be done for demonstrating significant gains under that assumption would motivate a case for enabling that functionality.

B.1.2 Inter-cell frame synchronization

In the last section, the problem of multiple symbol-types existing within the same VRB, was identified. And it was noted that this may serve as problematic in changing transmit powers on a per-VRB basis.

One possible solution to this is to coordinate interference in time rather than frequency, and coordinate the power among interfering cells, so that the concurrent neighbouring transmissions are time orthogonal rather than frequency orthogonal.

A question that naturally arises following such a thought is that of inter-cell frame synchronization. In the frequency coordinated case, many ICO schemes can be invented which require no frame synchronization. But obviously any ICO scheme that depends on temporal orthogonality of resources, would require synchronization to the level of accuracy associated with that orthogonality.

Fortunately, LTE will be deployed with full frame synchronization. The requirement set to the vendors is 4.8usec for macro cells. The current Femto cell deployments, which use the NTP and IEEE 1588v2 algorithms [9] are less strict but can achieve synchronization accuracies to the level of at least 10usec. LTE has a symbol length of 66.7usec and a subframe length of 1ms. These constraints imply that it would be possible to develop subframe-synchronized ICO algorithms, that would easily maintain appropriate temporal orthogonality.

B.1.3 Distributed versus centralized approaches

There are ICO algorithms which assume a centralized coordinator, for example [18, 6] and algorithms which assume distributed control, for example [3, 19]. It is interesting to ask what the feasibility of both methods are from a standardisation perspective.

In the centralized case, it is assumed that certain information is available to some network element above the base station. Typically it might be assumed that cell loads and SINR measures for each sub-band are available. In existing UMTS networks, certain statistics are already communicated to the RNC for the purpose of KPI monitoring and fault detection and thus it is feasible that in LTE a network node above the level of the eNB could be used to extract information and perform processing.

APPENDIX C

C.1 GLOSSARY

3GPP	Third Generation Partnership Project
ASB	Attenuated Sub Band
BS	Base Station
BSB	Boosted Sub Band
CQI	Channel Quality Indicator
DL	DownLink
ICIC	InterCell Interference Coordination
ICO	Interference Coordination
KPI	Key Performance Indicator
LTE	Long Term Evolution
MAC	Medium Access Control
MCS	Modulation and Coding Scheme
PRB	Physical Resource Block
QoS	Quality of Service
RSRP	Reference Signal Received Power
RSSI	Received Signal Strength Indicator
SINR	Signal to (Interference plus Noise) Ratio
TX	Transmitter
UE	User Equipment
UL	UpLink

APPENDIX D

D.1 Statistical tests under MATLAB

Assume that **a** is a 300x1vector containing the IQs of pupils aged 15 from school A and **b** is a 300x1 vector containing the IQs of pupils aged 15 from school B. It is postulated that 15 year old pupils from school A have higher IQs than those from school B because of socio-economic factors. We'll assume that experimental sampling has been carefully controlled so that any differences observed can be attributed to this and not something else.

In Matlab, to perform the test that pupils from school A are more intelligent than pupils from school B, the code is as follows:

```
[h,p] = ttest2(a,b,alpha,'right','unequal');
```

This performs a 'right' tailed test to decide whether the mean of **a** is greater than the mean of **b**, under the assumption that the vectors **a** and **b** are normally distributed, against the null hypothesis. The null hypothesis in this case is that the **a** and **b** come from distributions with equal means or that the mean of **b** is greater than the mean of **a**.

The value *alpha* signifies the probability at which the null-hypothesis should be rejected and is used to set the return value h which has the value h=1 if the null-hypothesis is rejected, and h=0 otherwise. The value *alpha* is usually set to 0.05, which means that h=1 when the probability that the null-hypothesis is true is less than or equal to 0.05 (5%).

The value 'unequal' is used to tell Matlab that we cannot assume that the variances of the two groups are equal. It is best to use this approximation whenever one is unsure about the variances of the two groups or when it is known they are unequal.

The return value p is the most important outcome, since h is simply set based on whether p is less than *alpha.*

The actual value of p gives the probability that the null hypothesis is correct. For example a value of 0.03 means that there is a 3% chance the null hypothesis is correct.

If one only cares whether **a** and **b** are different, then the test is:

```
[h,p] = ttest2(a,b,alpha,'both','unequal');
```

And to test whether the mean of **b** is greater than the mean of **a**, use:

```
[h,p] = ttest2(a,b,alpha,'left','unequal');
```

You might wonder about the assumption that **a** and **b** are normally distributed. Since the test relies on this assumption being true, for added assurance the correctness of this assertion can be verified using either a one-sample Kolmogorov-Smirnov test (in matlab [h,p] = kstest(**x**)), or Lilliefors test (in matlab [h,p] = lillietest(**x**)) where **x** is the input vector, for example **a** or **b** in the above example.

Chapter 5

Intrasite Scheduling for Interference Avoidance

Group Research & Development

C2 – COMPANY CONFIDENTIAL

Intrasite Scheduling for Interference Avoidance

Reference	VFGR&D_2010_0109
Version	1.0
Date	04 November 2010
Author	Ashley Mills
Theme/Function	Radio
Filed As	R&D Sharepoint https://sharepoint.vodafone.com/global/GroupRnD
Document type	Technical report
Status	Approved
Approved	David Lister
Signature / Date	.... / ....

Scope

This report addresses the choice of schedulers and their impact on interference mitigation techniques related to LTE.

3G/UMTS is out of scope.

Document History

Version	Date	Reason
0.5	22 October 2010	Draft
1.0	04 November 2010	Final

Document Distribution

Siavash Alamouti
Luke Ibbetson
Trevor Gill
David Lister
Peter Cosimini
Eric Murray
Ralf Irmer
Neil Scully
John Turk

Executive Summary

LTE will be deployed in a reuse one configuration, in which all frequency resources are available to use in each cell. This has generated concern over cell-edge performance. Numerous so-called interference mitigation techniques have been proposed to improve cell-edge performance. These vary from static deployments that require no communication between cells, to dynamic schemes which require frequent communication between widely distributed cells.

The former schemes have little merit in a hetrogeneous network, and the latter schemes, whilst promising, often assume inter-cell communication beyond the limits of the current technology.

This work examines an inteference avoidance strategy that should be feasible in LTE today, with only marginal modification of extant equipment. The work looks at communication within sectors of a given site only, and proposes a novel interference mitigation scheme within this framework.

The proposed scheme seeks a balance between intra-site resource orthogonality and independent frequency selectivity.

The method is compared against existing time and frequency selective scheduling approaches, as well as a static priority based interference avoidance scheme. The intra-site method promoted obtains better total and percentile throughputs compared to the competing schemes, and demonstrates the benefits that access to intra-site scheduling information can realise. Statistically robust results are obtained through system level simulations of LTE in the downlink direction only.

The gains obtained are relatively small, between 2% and 5% improvements in cell-edge and total cell throughputs, but are projected to come at little extra cost per site. The strategic implication threfore is a recommendation to encourage our equipment vendors to pursue intra-site scheduling for interference mitigation.

Future work will extend the proposed scheme with a view to improving the performance margin further.

Contents

1 INTRODUCTION

1.1 Motivation For Inter-Cell Interference Coordination

The next generation wireless technology, Long Term Evolution (LTE), has been designed to deliver higher spectral efficiency and increased cell-edge throughputs relative to HSPA [1]. It is expected that LTE will be deployed in a reuse one configuration, in which all frequency resources are available to use in each cell. Although LTE can operate at SINRs as low as -6.5dB [2], concern still persists over cell-edge performance.

This has led to the proposal of numerous inter-cell interference coordination mechanisms. A large group of these are static or semi static in nature [3, 4, 5, 6, 7]. However, upon examining results derived from these ideas [3, 8, 6, 9, 10, 11, 12, 13, 14] it is clear that while particular scenarios may benefit from one or another technique, for heterogeneous networks there is unlikely to be a net gain from static approaches [15].

Dynamic approaches which adapt to network conditions, and usually assume communication between basestations have thus been developed [16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 8, 27]. These schemes have tended toward taking more and more cells into account, and it would appear that the industry is converging toward multi-cell processing with a centralised RAN architecture [28, 29].

This work falls into the latter group. We present a novel means of interference avoidance for communicating groups of cells (illustrated for the intra-site case) which we believe could be applied now to real networks within the existing LTE standards framework. The method is compared with existing scheduling approaches which do not employ communication between cells.

The rest of this document is organised as follows: Section 1.2 provides some background concerning the different approaches toward inteference coordination and explains which approach the current work takes. Section 1.3 examines the potential for gains in intrasite coordination of resources through some simple monte-carlo experiments. Experimental methodology is provided in Section 2, with Section 0 giving an overview of the simulation approach taken. Sections 2.2 to Section 2.5 provide definitions of the interference avoidance and scheduling strategies tested. 2.1 specifies some LTE system level assumptions. The results are presented in Section 3, discussed in Section 4, and the work concludes in Section 5.

1.2 Differentiating between different interference scenarios

Figure 1 shows the throughput over 24h for two three-sector sites randomly selected from the Vodafone network in the UK. The data obtained from the network is in MB transferred per 15 minutes. The data presented has been converted into Mbps, binned as 15 minute averages.

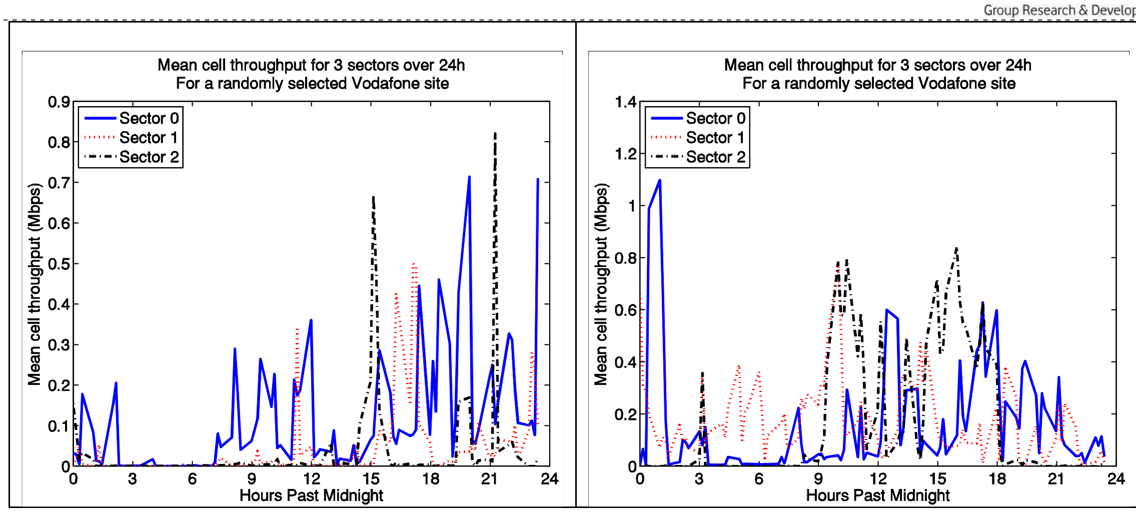

Figure 1: Traffic over a 24h period for the sectors of two randomly selected Vodafone basestations

Looking at the LHS, cell utilisation can qualitatively be described as low, in that the UMTS cells examined have a technical peak throughput in excess of 5Mbps yet nothing greater than approximately 0.8Mbps is seen. And although it is likely that the the instantaneous rates are much higher, since the rates plotted are 15 minute averages, the pertinent point is that it is unlikely to be the case that cells are operating in a full buffer state all the time.

Another artifact worth observing is that of the difference seen between sectors, and in this instance, between sectors of the same site. For example, in the LHS of Figure-1, Sector 0 seems generally to be more active than the others, but there are periods where Sector 1 and Sector 2 assume the highest activity. In the RHS of Figure-1 all sectors have similar levels of activities, but there are lots of temporal differences which set them apart.

The point of the last observation is to make it clear that interference from a given neighbour is unlikely to be constant, and will wax and wane in concert with the state of neighbour's data buffers. Conversely a cell is unlikely to be constantly interfering with other cells for the same reason.

Returning to LTE, now consider the sectors at a single site. From the perspective of any of the sectors, experienced inteference is a function of the bandwidth utilisation of the other two sectors at the time instant of signal reception. Interference will be experienced on any sub-band that the neighbours are using. The number of subbands experiencing interference can vary from 0 upto the channel limit. Finally, the interference can be coming from one neighbour, or both.

Given this information, and understanding that the aim is to coordinate bandwidth usage and thus interference, it should be clear that there are three principal scenarios:

1. No neighbours are interfering (they are not transmitting).
2. At least one neighbour is interfering on a *subset* of the available bandwidth.
3. At least one neighbour is interfering over *all* the bandwidth.

This limits the type of coordination strategies to the following:

1. No need for coodination.

2. Complete or partial orthogonality of allocations is possible.

3. Coordination with the fully utilized neighbour can only proceed by sacrificing its bandwidth for orthgonality.

The third coordination strategy would likely employ the technique of examining pairs of UEs to see whether they would benefit from pairwise orthogonal use of a resource blocks at the cost of reduced individual bandwidth.

This work begins by looking at the second type of coordination, where flexibility arises because bandwidth utilisation is less than 1 for neighbouring sectors. The approach taken will be to try and ensure that neighbours use ``good'' sets of orthogonal resource subbands wherever possible.

1.3 Potential for gains from intra-site scheduling

Before embarking on a course of deep and computationally expensive research into the gains that intra-cell scheduling can realize, it was considered prudent to perform a preliminary analysis of the gain potential.

In order to assess this potential, a two tier, 57 cell network was established, and the central cell populated with UEs at a resolution of one per $5m^2$. The LTE system assumptions were as in Section 2.1. To each UE, a single packet was sent from the eNB to obtain a value for the SINR at its location. Then, neighbour cells were switched off in the order shown in Figure 2, another packet sent, and the SINR re-evaluated for each case. The procedure was repeated for 100 random shadowing seeds and the results averaged.

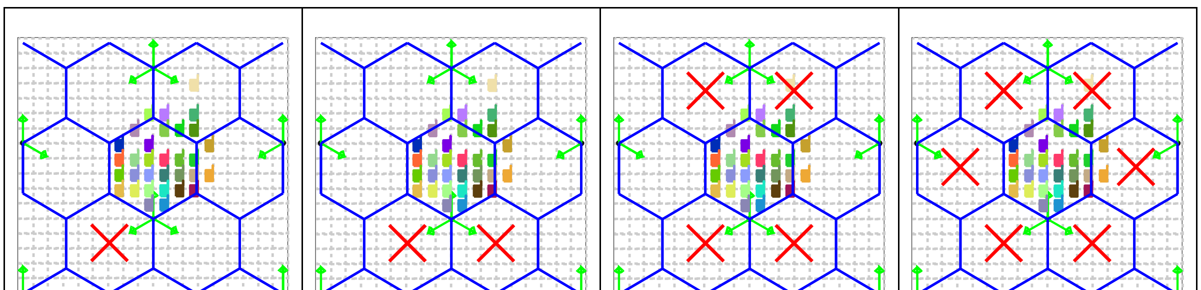

Set #A	***Set #B***	***Set #C***	***Set #D***

Figure 2: Each subfigure shows the central seven cells in a 57 cell two-tier hexagonal network. In the simulation, only the central cell was populated with UEs. The immediate neighbours were silenced in the manner illustrated in each set. What is not shown is the partner to *Set #A*: where the cell to the immediate right of the marked cell was instead silenced. The results for *Set #A* are an average over both cases. Note that for illustration purposes the grid resolution is drawn much larger than the value used in the simulations.

Figure 3 **(a)** plots the mean and std SINR improvement as increasing numbers of adjacent cells are silenced. In the case of ***Set #A*** the results were averaged over silencing either of the two immediate neighbours, to give an average for the ``single neigbor silenced'' case

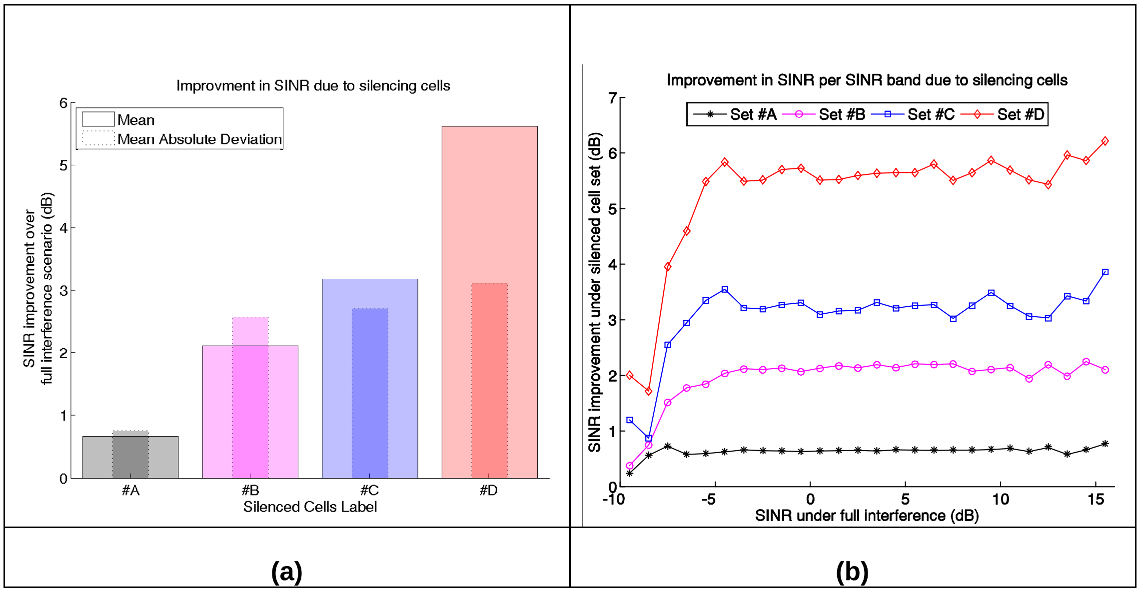

Figure 3: (a) Improvement in SINR attributed to silencing neighbours. (b) The same improvement but binned according to SINR values under no silencing.

As can be seen, silencing a single neighbour, on average obtains $\approx 0.66dB$ gain in SINR, with a mean absolute deviation of $\approx 0.75dB$. Silencing both intra-site neighbours more than doubles the mean improvement in SINR, giving $\approx 2.11dB$ gain on average, with $\approx 2.57dB$ mean absolute deviation. Silencing even more cells improves the SINR further as illustrated. Note that the mean absolute deviation is used in substitute of standard deviation because the gains seen are always positive and do not follow a normal distribution.

Figure 3 (**b**) shows the SINR improvement, binned according to the SINR under full interference. In each case, the improvement is lowest for very low SINRs, and improves with increasing SINR upto $\approx -5dB$ at which point it stays approximately constant for the remainder of the SINR range.

Note that the SINR under full interference is only plotted upto $16.5dB$. This is because a fully interfered hexagonal network, with the antenna pattern used here [2], and an intra-site correlation coefficient of one, has an SINR ceiling of $17dB$. When both intra-site interferes are removed, the SINR can be very high for locations close to the serving base station. This distorts the plot becuse the bin between $16.5dB$ and $17dB$ can have upto 10 times greater SINR improvements than the rest of the cell.

The effect is likely to occur in real life, but is likely to be softened by the non-ideal vertical radiation pattern which was not modeled here. It is also unclear whether an intrasite correlation of 1 is realistic since no measurement data was available. The effect affects less than 1/8th of 1% of the UEs examined. Note that the effect is not removed from the dataset in the dynamic simulations which follow, it was simply removed from Figure-3 for the sake of clarity. The beginnings of the effect can still be seen by the rise in SINR improvement in Figure 3 **(b)** for sets ***#C*** and ***#D***.

Looking at the SINR improvement associated with silencing both neighbours, it might seem intuitively strange that the gain is a lot more than twice the gain obtained by silencing only one neighbour. However, it turns out that the gain in the former case is *always* more than twice the latter, and only approaches exactly twice when the interference due to a single neighbour is negligible relative to sum of all interference. This can be seen by performing some simple algebra.

Let s represent the signal received by some user, and assume for simplicity that the interference received from each immediate neighbour cell (the intra-site neighbours) is equal to c, and that the interference received from the non-immediate neighbours, plus noise, is e. The signal to noise ratio experienced by the user can be written:

$$SINR_{full} = \frac{s}{e+2c}$$

The SINR in the case where one immediate neighbour cell is silenced, can be written:

$$SINR_{one} = \frac{s}{e+c}$$

And the SINR in the case where both immediate neighbours are silenced, can be written:

$$SINR_{two} = \frac{s}{e}$$

The improvement in SINR due to silencing one immediate neighbour can be written:

$$Improve_{one} = SINR_{one} - SINR_{full}$$

And the improvement due to silencing both immediate neighbours can be written:

$$Improve_{two} = SINR_{two} - SINR_{full}$$

The intuitive surprise is that $Improve_{two}$ is much greater than twice $Improve_{one}$. Now let us proceed to derive the factor $Improve_{two} \div Improve_{one}$:

$$Improve_{two} = \frac{s}{e} - \frac{s}{e+2c}$$

$$= \frac{s(e+2c)}{e(e+2c)} - \frac{se}{e(e+2c)}$$

$$= \frac{se+2sc-se}{e(e+2c)}$$

$$= \frac{2sc}{e(e+2c)}$$

$$Improve_{one} = \frac{s}{e+c} - \frac{s}{e+2c}$$

$$= \frac{s(e+2c)}{(e+c)(e+2c)} - \frac{s(e+c)}{(e+c)(e+2c)}$$

$$= \frac{se+2sc-se-sc}{(e+c)(e+2c)}$$

$$= \frac{sc}{(e+c)(e+2c)}$$

$$Improve_{two} \div Improve_{one} = Improve_{two} \cdot Improve_{one}^{-1}$$

$$= \frac{2sc}{e(e+2c)} \cdot \frac{(e+c)(e+2c)}{sc}$$

$$= \frac{2(e+c)sc(e+2c)}{esc(e+2c)}$$

$$= \frac{2(e+c)}{e}$$

And since e and c are always greater than zero (in the situation being examined), the SINR improvement due to silencing both immediate neighbours will always exceed twice the SINR improvement due to one (under the assumption the inteference is the same). Some modification is required to demonstrate this for the general case where the neighbour interferences differ, but the above should suffice to destroy any intuitive bias.

What this also suggests, is that the magnitude of e in the setup that produced Figure-3 is roughly proportional to the magnitude of c, since the increase in interference when both cells are silenced is roughly 4 times that of when only one cell is.

We have seen in this section that a mean SINR improvement of $\approx 2.1dB$ can be achieved by silencing neighbour cells. Note that silencing neighbours is equivalent to coordinating interference through orthgonal use of resources. A mean of about $0.5dB$ can be achieved by coordinating with only one neighbour. This clearly indicates that intra-site interference coordination has the potential to produce significant gains in throughput and justifies the investment in time required to continue with this thread of research.

2 METHODLOGY

2.1 LTE System Assumptions

Figure 4 illustrates the essential components of a DL LTE frame (in 10Mhz). The frame itself consists of 10 subframes which each last 1ms. Half of a subframe is called a slot. In frequency, each subframe is split into 50 Virtual Resource Blocks (VRBs). To the right of the frame a VRB is shown expanded.

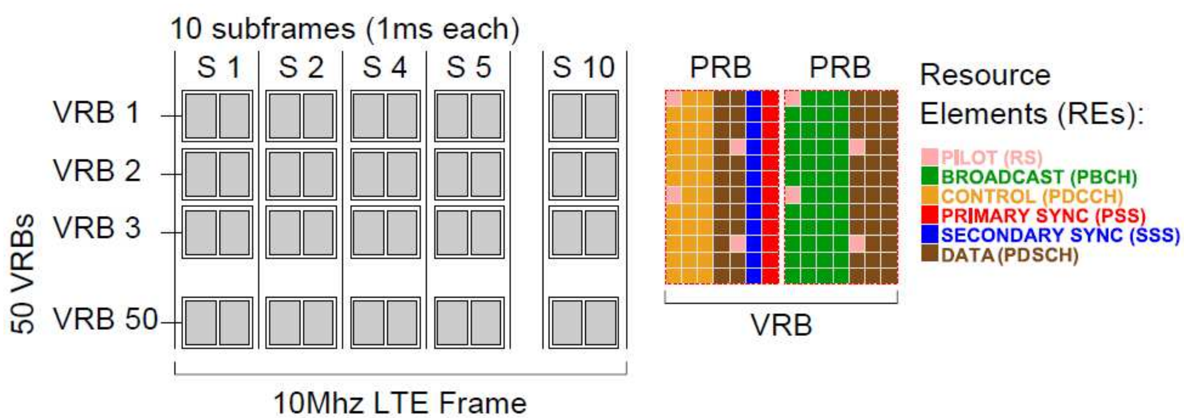

Figure 4: The essential components of a DL LTE frame (10Mhz Bandwidth)

Each VRB is comprised of a pair of Physical Resource Blocks (PRBs). Section 6.2.3 of 3GPP 36.211 [30] states that PRBs are mapped to VRBs under a distributed or localized mapping, and are assigned to UEs in pairs (one from each slot of the subframe) under a single VRB number. Accordingly this text uses the VRB as an atomic unit of allocation, since it is in fact technically impossible to independently allocate the PRBs within a VRB. Each PRB of a VRB spans 12 subcarriers in frequency and 7 symbols in time (note that for large inter-site distances a longer cyclic prefix is available which spans only 6 symbols in time, but that here the former duration is used). Each unit in a PRB is called a Resource Element (RE). An RE spans a single subcarrier in frequency and single symbol in time. An RE has a frequency width of 15khz and lasts approximately 70µs.

An RE encodes exactly one MCS symbol. For this reason the term symbol is often used interchangeably with RE. LTE uses REs for a number of different purposes in addition to serving data. The RE types modelled here are shown, with their locations in the VRB to the right of Figure-5. The following list describes their functions.

- Control channels

  Control channels are used here only to model VRB RE usage. Decoding of control channels and associated decoding failures are not modelled. Control channels are assumed to take up the first three RE columns of each VRB.

- Pilot symbols

Pilot symbol placement is modelled accurately according to [30]. The computation of RSRP from pilot symbols is performed as defined in Section 5.11 of [31]. RSRP measurements are used to associate UEs with serving BSs in the static mobility simulations. The figure shows the pilot symbol positions for a single spatial channel on a cell having CELL_ID = 0.

- Synchronization and broadcast channels.

Synchronization and broadcast channels are included only to model VRB RE usage accurately. Primary and secondary synchronization symbols are only present on the 3 VRB either side of the central carrier and are only present in subframes 0 and 5. Broadcast channels are only present on the 3 VRB either side of the central carrier and are only present in subframe 0. They are shown in the VRB diagram only to indicate their positions when they do occur, but be aware that most of the time data REs take their place.

As already stated, synchronization and broadcast channel mechanisms are not simulated, they only consume space, so as to model the RE usage accurately.

- Data symbols

Since only subframes 0 and 5 contain synchronization and broadcast channels, the number of data sresource elements per VRB, changes depending on the subframe and location within the band. Taking this into consideration, an average number of data REs per VRB can be computed.

Factoring in the presence of control and pilot symbols, the average number of data resource elements per VRB for the simulator is 124.8720 data symbols.

Overview

Four scheduling schemes were compared: one novel intrasite scheme (FD-PF-Intrasite) and three existing non-intrasite schemes (FD-PF, TD-PF-RD, TD-PF-SS). The schemes are defined in the subsections which follow. Each scheduling strategy was tested for 9 different parameter configurations: the Cartesian product of three cell load conditions U and three scheduling biases α.

U represents the traffic load in each cell. $U = x$ means that each sector is limited to assigning exactly the fraction x of all available VRB within one TTI. The motivation for this approach was that it was desired to abstract from a particular traffic class and capture the essence of intra-site interference. From the perspective of a serving cell at a given site only three scenarios are relevant:

1. Low interference/load: each sector uses less than 1/3rd bandwidth. Full resource orthogonality among sectors is in principle possible.

2. Partial interference/load: each cell uses over 1/3rd of the bandwidth, but less than the full bandwidth. Partial orthogonality among sectors is possible.

3. Full interference/load: no orthogonality is possible among sectors.

These three scenarios capture the possible types of intra-site interference that a sector can experience, independent of traffic class. The above scenarios are represented here by U values of 0.3, 0.65, and 1 respectively.

α is defined in Equation 1 and controls how much emphasis is put on the instant rate exponent in the proportional fair weighting of UEs. For $\alpha > 1$, the scheduler approaches a Maximum C/I scheduler and for $\alpha < 1$, the scheduler approaches a round robin scheduler, for $\alpha = 1$ a normal PF approach is followed. The values 0.1, 1, and 10 were examined here.

For each of the 9 parameter tuples, and for each scheduling strategy, a system simulation was executed to collect throughput statistics. Each simulation was performed on a two-tier network of 57 cells with wraparound and a lognormal shadowing standard deviation of $8dB$. The FD-PF-Intrasite and site to site correlations assumed values of 1 and 0.5 respectively. Each simulation was repeated for 50 different shadowing seeds, and the collated results averaged.

Each simulation lasted for a duration of 1.25 seconds, the first 0.25 seconds of which were used for warmup during which round robin scheduling was used and no statistics were collected. Fast fading was modelled using the SCME sub-urban macro model [32] assuming a constant UE velocity of $1.5m/s$.

A 3dB offset Shannon mapping between SINR and bitrate was employed with a ceiling of 5.52 bits per symbol. This corresponds to 64QAM at a coderate of 0.92, which is the maximum available. MCS targets were assumed to exist at 1dB intervals, and packets received with 0.1dB or less than the target SINR were assumed to have been received in error. Note that this approximation produces a curve very similar to realistic throughputs seen by trial equipment, but owing to the simplifications inherent in the model, is only suitable for relative comparisons such as this work.

Parameter		Value	Unit
Number of Cell Sites		57	
Number of Sectors Per Site		3	
Inter-Site Distance		500	metres
Lognormal Shadowing Standard Deviation		7	dB
Shadowing Correlation	Between Cells	0.5	
	Between Sectors	1.0	
BS Antenna Boresight Gain		14	dBi
BS Horizontal Antenna Pattern		$A(\theta) = -\min\left[12\left(\frac{\theta}{\theta_{3dB}}\right)^2, A_m\right]$ θ_{3dB} = 70 degrees, A_m = 20 dB	dB
Carrier Frequency		2000	MHz
System Bandwidth		600 subcarriers x 15 kHz = 9000	kHz
			[33]
Base-Station Antenna Height		30	metres
UE Antenna Height		1.5	metres
UE Noise Figure		9	dB
Minimum Coupling Loss		60	dB

Figure 5: System level simulation assumptions

2.2 FD-PF (Frequency Domain Proportional Fair)

Time domain proportional fair [34] has been extended to the frequency domain in various ways and gives better performance than the former (See for example [35, 36, 37, 38]). The approach used here to determine the score $S_{i,m}(t)$ for UE i on subband m at time t is as follows:

$$S_{i,m}(t) = R_{i,m}(t)^{\alpha} / MR_i(t)^{\beta} \quad (1)$$

Where $R_{i,m}(t)$ is the instantaneous rate of UE i on subband m (in bits per symbol) at time t, $MR_i(t)$ is the mean rate of UE i at time t. The mean rate is updated recursively as follows:

$$MR_i(t+1) = A_i(t)\cdot(1-1/t_c) + MR_i(t)\cdot 1/t_c \qquad (2)$$

Where $A_i(t)$ is the payload allocated to UE i at time t. In this work, payload was expressed simply in bits. The parameter t_c controls the size of the averaging window, assuming a value of 0.01. Note that only one of α and β is needed in practise.

Exactly 9 subbands were employed here containing 6 VRB each in accordance with the LTE specification [39] for $10Mhz$ (See Table 7.2.1-3). At TTI t, subband m is assigned to the UE i^* with highest score $S_{i^*,m}(t) = \underset{i}{argmax}\, S_{i,m}(t)$.

2.3 TD-PF-RD (Time Domain Proportional Fair, Random Distributed)

Time domain proportional fair is a special case of frequency domain proportional fair where the number of subbands is set to 1. The random distributed aspect refers to the ordering of VRBs during assignment. Assignment begins at VRB index k, which is chosen randomly, and proceeds contiguously, wrapping back to index 0 if necessary. The effect is that interference is partially distributed among neighbours.

2.4 TD-PF-SS (Time Domain Proportional Fair, Semi Static)

The TD-PF-SS approach differs from the TD-PF-RD approach described above, only in the order in which VRB are allocated within each cell. Instead of starting at a random VRB index, the start indices are coordinated in a tessellating pattern with adjacent cells in a manner analogous with reuse three. For three sectors at a given site with $10Mhz$ bandwidth: sector 0 starts scheduling at VRB index 0, sector 1 starts at VRB index 16, and sector 2 starts at VRB index 32. This pattern tessellates in hexagons, such that if all sectors are using less than 1/3rd of the bandwidth, then their VRB allocations will be orthogonal. Note that the TD-PF-SS approach described is essentially identical to that proposed in [40].

2.5 FD-PF-Intrasite (Frequency Domain Proportional Fair, FD-PF-Intrasite Variant)

This approach is novel, and is inspired by the TD-PF-SS approach, and the observation that the latter, although being able to obtain resource orthogonality for low loads, is reduced in frequency selectivity due to the consequent restrictions on VRB use. The motivation behind this approach is to capitalise on the benefits of both resource orthogonality and frequency selectivity. The resource assignment algorithm, executed *for each site*, is shown in Procedure 1 below.

Construct a three dimensional scheduling ``matrix" having:

- 150 rows: one row for each VRB at the site such that each consecutive 50 entries corresponds to a different sector.
- For each row, make as many columns as there are scheduled UEs at the corresponding sector.

Populate and order the scheduling matrix:

1. Obtain the scheduling score for each row (site-level VRB index) and column (UE) using the FD-PF scoring metric.
2. Sort each row by column according to UE scores.
3. Sort the rows by the UE score at the head of each row.

Schedule the UEs:

- Iterate three times through the matrix by row:
  - (a) Determine the sector-level VRB index for the row.
  - (b) Determine the parent sector for the current row.
  - (c) *IF* a VRB with this sector-level index has not been allocated this iteration *AND* the VRB has not already been allocated to the parent sector *THEN*:
    - Allocate the VRB to the UE at the head of the current row.

Procedure 1: Procedure for intra-site scheduling.

The algorithm begins by visiting each sector at the site independently, to obtain per sector, per VRB scores for each UE (according to the FD-PF algorithm). It then creates a site-level meta-score for each VRB using the highest scoring UE for that sector and on that VRB. Finally the site-level scores are sorted so that for a given VRB index the best scoring UE among all sectors at the site gets priority. A three round iterative process ensures that within a given round, a VRB with a given index can only be assigned to one sector. This ensures orthogonal assignment of VRB among sectors during each round, such that if the utilisation in each sector is less than 1/3rd, then VRBs assigned to each sector are completely orthogonal. It also ensures however, that when all neighbours are fully utilised, the algorithm will behave the same as the frequency selective scheduler.

In essence the aim of the algorithm is to find good orthogonal sets of VRBs for low utilisation scenarios, without sacrificing performance at high neighbour utilisations.

3 RESULTS

3.1 Total Throughputs

Table 1 compares the results of the four scheduling strategies for each proportional fair instant rate exponent α, and for each utilisation value U.

α	*U*	TD-PF-SS	TD-PF-RD	FD-PF	FD-PF-Intrasite
0.1	0.3	1.96	0.54	2.71	**2.74**
1	0.3	4.55	1.13	4.51	**4.74**
10	0.3	**9.02**	3.38	8.51	8.76
0.1	0.65	2.98	1.85	4.05	**4.20**
1	0.65	6.71	3.83	7.62	**8.02**
10	0.65	**16.23**	10.73	15.44	15.73
0.1	1.0	3.86	3.86	**5.37**	**5.37**
1	1.0	8.57	8.57	**11.34**	**11.34**
10	1.0	21.17	21.17	**21.50**	**21.50**

Table 1: System level results. α represents the exponent of the instantaneous rate of the proportional fair priority metric, and U represents the cell utilisation factor. As explained previously. For each row, the highest score is highlighted in bold font.

Within a given row, the four scheduling columns are comparable. The results are grouped into three sets according to the utilisation factor U.

Consider the performance of the FD-PF-Intrasite scheduler relative to the TD-PF-SS scheduler. According to a right tailed ttest with a significance threshold of 0.05, with Satterthwaites approximation to address the Behrens-Fisher problem [41], the FD-PF-Intrasite scheduler, FD-PF-Intrasite, significantly outperforms the semi static TD-PF-SS scheduler in all but two cases. The two cases are for a PF instantaneous rate exponent of 10, where U is 0.3 and 0.65 respectively. This is because in these cases the scheduling is heavily biased toward the highest SINR UEs and the resources under the TD-PF-SS scheduler between interfering neighbours are orthogonal or near-orthogonal respectively.

The FD-PF-Intrasite scheduler significantly outperforms the FD-PF scheduler for every 0.3 and 0.65 load point, except when $\alpha = 0.1$ and $U = 0.3$. In the latter case the p value[1] is 0.07 however, so is close to being significant at a threshold of 0.05.

For $U = 1.0$, the FD-PF-Intrasite and FD-PF schedulers have the same performance, because the algorithms behave exactly the same at full load. Similarly the TD-PF-SS and TD-PF-RD algorithms have the same performance for $U = 1.0$.

The FD-PF-Intrasite scheduler significantly outperforms the TD-PF-RD scheduler for every point.

The improvement of the FD-PF-Intrasite scheduler over the FD-PF scheduler is maximum at $\approx 5\%$ (as a percentage of the FD-PF value), and has a mean of $\approx 2.2\%$ with a standard deviation of $\approx 2.1\%$.

The improvement of the FD-PF-Intrasite scheduler over the the TD-PF-SS scheduler is maximum at $\approx 40\%$ (as a percentage of the TD-PF-SS value), and has a mean of $\approx 19\%$ with a standard deviation of $\approx 19\%$.

Finally, it is worth noting the difference between frequency *and* time selectivity compared to *only* time selectivity. The improvement of FD-PF over TD-PF-RD is maximal at 500% (as a percentage of the TD-PF-RD value), and has a mean of $\approx 232\%$ with a standard deviation of $\approx 36\%$.

---

[1] The p value represents the probability that the observed differences are due to chance. A p value of 0.07 means that there is a 7% chance that the observed differences are due to chance, under the assumption that the distributions examined are normal.

3.2 Throughput Distribution

The CDF of throughput is interesting to examine in the context of the total throughput results, since it indicates to whom that throughput is distributed. Upon examining the CDFs for each of the input configurations examined in Table 1 it is however apparent that the general behaviour is very similar across these configurations. Therefore, in the interests of brevity, only three CDFs will be examined here. Figure 6 shows the throughput CDF for $U = 0.65$ and $\alpha = 0.1$.

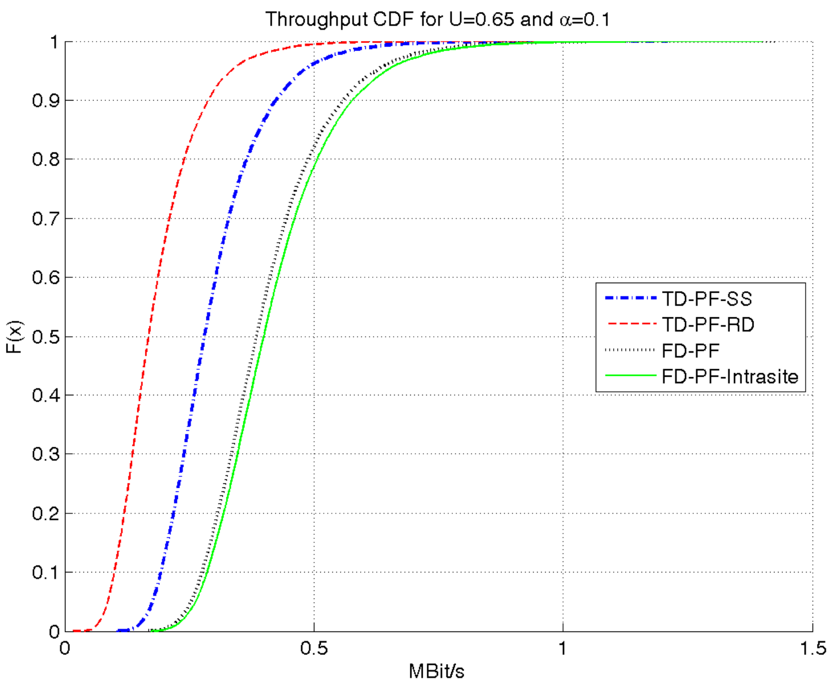

Figure 6: Throughput CDF for $U = 0.65$ and $\alpha = 0.1$

As can be seen, the improvements of the FD-PF and FD-PF-Intrasite schemes over the TD schemes are uniform over the distribution, with perhaps a slight increase in improvement toward higher throughputs. The improvement of the FD-PF-Intrasite scheme over the FD-PF scheme also appears very uniform. Figure 7 shows the throughput CDF for $U = 1.0$ and $\alpha = 1$.

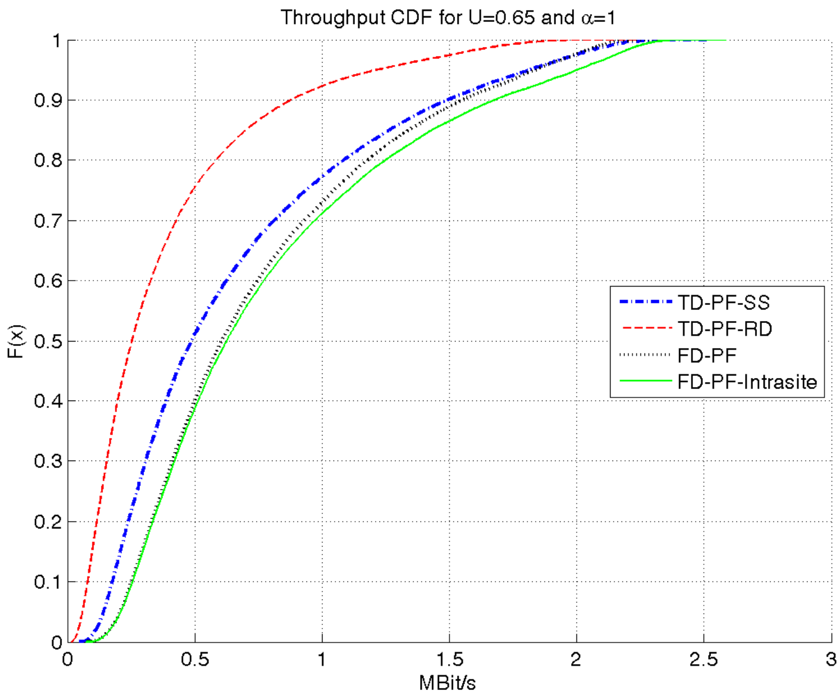

Figure 7: Throughput CDF for $U = 0.65$ and $\alpha = 1$

The impact of the α value can be seen by comparing both plots. For example, the percentage of UEs getting more than 0.5Mbit/s in Figure 6 under the FD-PF-Intrasite scheme is $\approx 20\%$ whereas in Figure 7, this percentage is increased to $\approx 60\%$. This should make it clear how much influence the scheduling strategy has on the tradeoff between percentile and total cell throughputs.

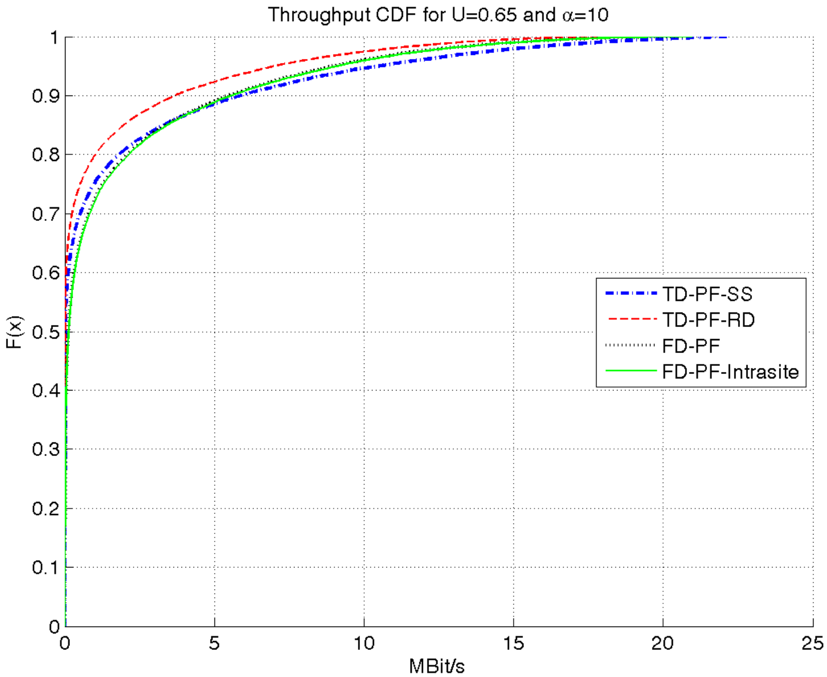

Figure 8: Throughput CDF for $U = 0.65$ and $\alpha = 10$

In Figure 7 the gain of the FD-PF-Intrasite scheme over the FD-PF scheme is greater for higher throughputs (that is, higher SINR UEs). The gain of the FD-PF-Intrasite and FD-PF schemes over the TD-PF-SS scheme diminishes toward the higher SINR UEs.

Figure 8 shows the CDF for $U = 0.65$ and $\alpha = 10$ and illustrates why the TD-PF-SS scheme does better than the FD-PF-Intrasite and FD-PF schemes for this point: it lies to the right of these schemes at the high end of the throughput distribution, and although it lies to the left at the lower end of the throughput distribution, in trade-off the result is that the TD-PF-SS scheme gets a higher throughput.

Notice however that with $\alpha = 10$ the scheduling strategy is so aggressive that it is approaching a Maximum C/I type scheduler. Closer inspection of the input data reveals that a large percentage of the UEs are never selected for scheduling under $\alpha = 10$. If the simulation were run for an infinite amount of time, the CDFs may lie marginally to the right, but the relative performance of schemes is unlikely to change. A more pertinent observation is simply that $\alpha = 10$ is an unrealistic scheduling strategy, albeit a useful illustration in this case.

4 DISCUSSION

When looked at over the full range of scenarios examined, it should be clear that FD-PF-Intrasite scheduling strategy gives the best total and percentile cell throughputs. Although the mean improvement in total throughput of FD-PF-Intrasite over FD-PF is only 2.2% it should be observed that this gain comes about at very little cost: the modification the eNB probably being a matter of a software change. Furthermore, the gain is statistically robust to changes in shadowing distribution, and therefore the actual gain is likely to vary from site to site, and of course varies with neighbour loading. The recommendation therefore is strongly that these schemes should be pursued for deployment in LTE.

An interesting observation is that differences in throughput due to the scheduler exponent α, are almost as large as the differences that can be attributed to the load parameter U. With this in mind it is important that when judging the performance of a system, that the interaction of scheduling and load parameters are taken into account. It should also be clear that specification of a PF scheduler without providing α is meaningless.

4.1 Future work

The FD-PF-Intrasite scheme tested here provides no benefit over FD-PF for $U = 1$, since in this case the schemes perform identically. It has always been the intention to extend the scheme so to obtain differentiation in the case of $U = 1$, and this will be the subject of future work. In brief, the investigation will proceed in two directions:

1. Pairwise orthgonality enforcement

   Pairs of UEs, or in principle triplets of UEs in neighbouring cells will be put forth as candidates for enforced orthogonality at the VRB level. This is likely to produce a net gain for low SINR UEs, as evidenced by our previous work on static interference schemes [15].

2. Improved CQI estimation

   The discrepancy between neighbour activity at the point of CQI measurement and the point of CQI use, will be used to improve the accuracy of the CQI information. This should result in increased BLER and more appropriate matching of MCS to UE. The overall outcome should be improved cell throughputs.

5 CONCLUSION

The FD-PF-Intrasite scheduling scheme has shown that robust gains can be obtained over neighbour-naive frequency selective scheduling. The vendors should be encouraged in developing such strategies.

References

[1]3GPP TR 25.913, V9.0.0: Requirements for Evolved UTRA (E-UTRA) and Evolved UTRAN (E-UTRAN), 2009.

[2]3GPP TS 36.942 - V8.2.0 - Evolved Universal Terrestrial Radio Access (E-UTRA); Radio Frequency (RF) system scenarios, May 2009.

[3]Huawei. 3GPP TSG RAN WG1 Meeting #41, R1-050507 - Soft Frequency Reuse Scheme for UTRAN LTE, 2005.

[4]Siemens. 3GPP TSG RAN WG1 Meeting #41, R1-050476: Evolved UTRA uplink scheduling and frequency reuse. 2005.

[5]Ericsson. 3GPP TSG RAN WG1 Meeting #42, R1-050764 - Inter-cell Interference Handling for E-UTRA, 2005.

[6]Alcatel. 3GPP TSG RAN WG1 Adhoc Meeting, R1-050593 - Interference coordination for evolved UTRA uplink access, 2005.

[7]Texas Instruments. R1-051059: Inter-Cell Interference Mitigation for EUTRA, October 2005.

[8]F. Xiangning, C. Si, and Z. Xiaodong. An Inter-Cell Interference Coordination Technique Based on Users Ratio and Multi-Level Frequency Allocations. In *Wireless Communications, Networking and Mobile Computing (WiCom 2007)*, pages 799–802, September 2007.

[9]NTT DoCoMo Ericsson. TSG-RAN WG1 meeting #44, R1-060586 - Downlink and uplink inter-cell inteference co-ordination/avoidance - impact on the specifications, 2006.

[10]Mohammad Abaii et al. IST-4-027756. WINNER II D4.7.2 v1.0. Interference avoidance concepts.

[11]Texas Instruments. 3GPP TSG RAN WG1 Meeting #44, R1-060368 - Performance of Inter-Cell Inteference Mitigation with Semi-Static Frequency Planning for EUTRA Downlink.txt, 2006.

[12]Ericsson. R1-061374: Downlink inter-cell interference co-ordination/avoidance evaluation of frequency reuse, May 2006.

[13]IPWireless. 3GPP RAN WG1 Ad Hoc Meeting, R1-050652 - Attaining the Cell Edge Performance Requirements for the LTE Downlink, 2005.

[14]Siemens. 3GPP RAN WG1 Ad Hoc Meeting, R1-060135 - Interference Mitigation by Partial Frequency Reuse, 2006.

[15]Ashley Mills and David Lister and Marina De Vos. Understanding Static Inter-Cell Interference Coordination Mechanisms in LTE. *Vodafone Group R&D Internal Report: VFGR&D_2010_0045*.

[16]Mahmudur Rahman and Halim Yanikomeroglu and William Wong. Interference Avoidance with Dynamic Inter-Cell Coordination for Downlnink LTE System. In *WCNC'09: Proceedings of the 2009 IEEE conference on Wireless Communications & Networking Conference*, pages 1238–1243, Piscataway, NJ, USA, 2009. IEEE Press.

[17]Alexander L. Stolyar and Harish Viswanathan. Self-organizing Dynamic Fractional Frequency Reuse for Best-Effort Traffic Through Distributed Inter-cell Coordination. In *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, pages 691–699, April 2008.

[18]Weiwei Wu, M. Gitlits, and T. Sakurai. Dynamic resource allocation with inter-cell interference coordination for 3GPP LTE. *Microwave Conference, 2008. APMC 2008. Asia-Pacific*, pages 1 –4, dec. 2008.

[19]Zhifeng Tao et al and Toshiyuki Kuze. Dynamic Inter-cell Interference Coordination (ICIC) and Signaling, 2008.

[20]Xunyong Zhang, Chen He, Lingge Jiang, and Jing Xu. Inter-cell interference coordination based on softer frequency reuse in ofdma cellular systems. pages 270 –275, jun. 2008.

[21]R.Y. Chang, Zhifeng Tao, Jinyun Zhang, and C.-C.J. Kuo. Multicell ofdma downlink resource allocation using a graphic framework. *Vehicular Technology, IEEE Transactions on*, 58(7):3494 –3507, sep. 2009.

[22]N. Himayat, S. Talwar, A. Rao, and R. Soni. Interference management for 4g cellular standards [wimax/lte update]. *Communications Magazine, IEEE*, 48(8):86 –92, aug. 2010.

[23]Ericsson. R1-074444: On Inter-cell Coordination Schemes without/with Traffic Load Indication. *3GPP TSG-RAN WG1 Meeting #50, Shanghai, China, October 2007.*

[24]Marc C. Necker. Scheduling Constraints and Interference Graph Properties for Graph-based Interference Coordination in Cellular OFDMA Networks. 14(4):539–550, 2009.

[25]M.C. Necker. Towards Frequency Reuse 1 Cellular FDM/TDM Systems. In *Proceedings of the 9th ACM/IEEE International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM 2006)*, October 2006.

[26]Marc C. Necker. Integrated scheduling and interference coordination in cellular ofdma networks. pages 559 –566, sep. 2007.

[27]Gabor Fodor and Chrysostomos Koutsimanis and András Rácz and Norbert Reider and Arne Simonsson and Walter Muller. Intercell Interference Coordination in OFDMA Networks and in the 3GPP Long Term Evolution System. *Journal of Communications*, 4(7):445–453, August 2009.

[28]Gary Boudreau and John Panicker and Ning Guo and Rui Chang and Neng Wang and Sophie Vrzic. Interference Coordination and Cancellation for 4G Networks. *IEEE Communications Magazine*, 47(4), April 2009.

[29]China Mobile Research Institute. C-RAN: The Road Towards Green RAN. White Paper. Version 1.0.0., April 2010.

[30]3GPP TS 36.211 - V8.7.0 - Evolved Universal Terrestrial Radio Access (E-UTRA); Physical Channels and Modulation, May 2009.

[31]3GPP TS 36.214 - V8.7.0 - Evolved Universal Terrestrial Radio Access (E-UTRA) Physical layer – Measurements (Release 8), September 2009.

[32]Danian S. Baum, Jan Hansen, Giovanni Del Galdo, Marko Milojevic, Jari Salo, and Pekka Kyoesti. An Interim Channel Model for Beyond-3G Systems: Extending the 3GPP Spatial Channel Model (SCM). In *IEEE 61st Vehicular Technology Conference*, volume 5, pages 3132–3136, Stockholm, Sweden, May 2005.

[33]Digital Mobile Radio Towards Future Generation Systems. Cost 231 Final Report. Chapter 4: Propagation Prediction Models. http://www.lx.it.pt/cost231/docs/PDFfiles/Chap4_pdf.ZIP.

[34]Jalali, A. and Padovani, R. and Pankaj, R. Data throughput of CDMA-HDR a high efficiency-high data rate personal communication wireless system. In *Vehicular Technology Conference Proceedings, 2000. VTC 2000-Spring Tokyo. 2000 IEEE 51st*, volume 3, pages 1854–1858 vol.3, 2000.

[35]Kian Chung Beh and Simon Armour and Angela Doufexi. Joint Time-Frequency Domain Proportional Fair Scheduler with HARQ for 3GPP LTE Systems. In *Vehicular Technology Conference, 2008. VTC 2008-Fall. IEEE 68th*, pages 1 –5, Sep. 2008.

[36]Mongha, G. and Pedersen, K.I. and Kovacs, I.Z. and Mogensen, P.E. QoS Oriented Time and Frequency Domain Packet Schedulers for The UTRAN Long Term Evolution. In *Vehicular Technology Conference, 2008. VTC Spring 2008. IEEE*, pages 2532–2536, May 2008.

[37]Wang Anchun and Xiao Liang and Zhou Shidong and Xu Xibin and Yao Yan. Dynamic resource management in the fourth generation wireless systems. volume 2, pages 1095–1098, April 2003.

[38]Pokhariyal, A. and Pedersen, K.I. and Monghal, G. and Kovacs, I.Z. and Rosa, C. and Kolding, T.E. and Mogensen, P.E. HARQ Aware Frequency Domain Packet Scheduler with Different Degrees of Fairness for the UTRAN Long Term Evolution. In *Vehicular Technology Conference, 2007. VTC2007-Spring. IEEE 65th*, pages 2761–2765, April 2007.

[39]3GPP TS 36.213 V9.2.0 - Evolved Universal Terrestrial Radio Access (E-UTRA); Physical layer procedures (Release 9), June 2010.

[40]LG Electronics. 3GPP TSG RAN WG1 Meeting #42, R1-050833 - Interference mitigation in evolved UTRA/UTRAN, 2005.

[41]B D Hall and R Willink. Does “Welch-Satterthwaite” make a good uncertainty estimate? *Metrologica*, 38(1), 2001.

A.1 GLOSSARY

3GPP	Third Generation Partnership Project
ASB	Attenuated Sub Band
BS	Base Station
BSB	Boosted Sub Band
CQI	Channel Quality Indicator
DL	DownLink
ICIC	InterCell Interference Coordination
ICO	Interference Coordination
KPI	Key Performance Indicator
LTE	Long Term Evolution
MAC	Medium Access Control
MCS	Modulation and Coding Scheme
PRB	Physical Resource Block
QoS	Quality of Service
RSRP	Reference Signal Received Power
RSSI	Received Signal Strength Indicator
SINR	Signal to (Interference plus Noise) Ratio
TX	Transmitter
UE	User Equipment
UL	UpLink

APPENDIX B

B.1 Statistical tests under MATLAB

Assume that **a** is a 300x1vector containing the IQs of pupils aged 15 from school A and **b** is a 300x1 vector containing the IQs of pupils aged 15 from school B. It is postulated that 15 year old pupils from school A have higher IQs than those from school B because of socio-economic factors. We'll assume that experimental sampling has been carefully controlled so that any differences observed can be attributed to this and not something else.

In Matlab, to perform the test that pupils from school A are more intelligent than pupils from school B, the code is as follows:

```
[h,p] = ttest2(a,b,alpha,'right','unequal');
```

This performs a 'right' tailed test to decide whether the mean of **a** is greater than the mean of **b**, under the assumption that the vectors **a** and **b** are normally distributed, against the null hypothesis. The null hypothesis in this case is that the **a** and **b** come from distributions with equal means or that the mean of **b** is greater than the mean of **a**.

The value *alpha* signifies the probability at which the null-hypothesis should be rejected and is used to set the return value h which has the value h=1 if the null-hypothesis is rejected, and h=0 otherwise. The value *alpha* is usually set to 0.05, which means that h=1 when the probability that the null-hypothesis is true is less than or equal to 0.05 (5%).

The value 'unequal' is used to tell Matlab that we cannot assume that the variances of the two groups are equal. It is best to use this approximation whenever one is unsure about the variances of the two groups or when it is known they are unequal.

The return value p is the most important outcome, since h is simply set based on whether p is less than *alpha.*

The actual value of p gives the probability that the null hypothesis is correct. For example a value of 0.03 means that there is a 3% chance the null hypothesis is correct.

If one only cares whether **a** and **b** are different, then the test is:

```
[h,p] = ttest2(a,b,alpha,'both','unequal');
```

And to test whether the mean of **b** is greater than the mean of **a**, use:

```
[h,p] = ttest2(a,b,alpha,'left','unequal');
```

You might wonder about the assumption that **a** and **b** are normally distributed. Since the test relies on this assumption being true, for added assurance the correctness of this assertion can be verified using either a one-sample Kolmogorov-Smirnov test (in matlab [h,p] = kstest(**x**)), or Lilliefors test (in matlab [h,p] = lillietest(**x**)) where **x** is the input vector, for example **a** or **b** in the above example.

Chapter 6

Current Developments

This chapter discusses how the work presented in the previous chapter can be advanced, and provides some preliminary results to indicate the likely success of the suggested approach.

6.1 Leveraging Intrasite Gains

A clear avenue for continuation of the present work is to improve upon the algorithm presented in Chapter 5. The algorithm was designed to operate well at low to medium interference but operates (literally) identically to sector-independent frequency selective scheduling for the full interference scenario, where all cells in the network are simultaneously transmitting. This is because the algorithm tries to exploit the frequency and diversity gains associated with joint scheduling of three cells. Whereas in the case of full interference, *there is no additional frequency or diversity gain*, since all transmissions collide. Obtaining benefits in the face of full interference is a difficult problem. There are two general solutions:

1. Reduce the transmit power (possibly to zero) within a cell on some resource, to the benefit of another cell. The other cell may optionally increase its power in proportion to the reduction of the former.

2. Enforce orthogonality between n UEs by giving them a proportion $1/n$ of a resource of which they were a priori *all* receiving a proportion n (i.e. they were previously

receiving and colliding on all of the resource in question).

In the first case, the donor potentially loses out, whereas in the latter in principle any of the partners can lose out if the gains associated with the transaction are not equally distributed. And if orthogonality is enforced inappropriately, then all cells involved might lose out. The problem of inter-cell fairness is thus raised, as well as the need to successfully identify the conditions under which participants will profit from a change in resource assignment.

Given the work on soft frequency reuse in Chapter 4, the gains from downlink power control, in the form of soft reuse, appear to be limited to particular scenarios. Furthermore, the mutually exclusive use of resources tended to perform better than the soft sharing of resources. For these reasons, the proposal for continuation explored here will focus on the second item above; namely the enforcement of orthogonality between suitable candidates, at the cost of bandwidth.

The main problem in exploiting this idea, is that in the absence of inter-cell communication, a given cell does not know the resource assignments of its neighbours, at least not to any meaningful extent. To see why this is important, imagine that a UE makes a CQI measurement at a point in time when both its neighbours happen to be transmitting. Now imagine that at later time when the CQI report is used, i.e. at the point of transmission, that none of the immediate neighbours are going to transmit. In this case, there will be less interference from the neighbours when the packet is transmitted than there was when the CQI measurement was made. And since CQI measurements inform MCS selection, this could result in an inappropriately low MCS selection, and consequent under-utilisation of the channel. The converse problem occurs when the interference at the point of transmission is greater than expected.

This problem is shown to have a significant non-negligible effect on throughput and coverage in [123]. However, with an intrasite scheduler, or inter-cell communication, *it is* possible to know the resource assignments of the immediate neighbours. The discrepancies in measured CQI and experienced SINR can thus potentially be rectified.

Unfortunately, even if the resource assignments of neighbour cells *are known*, it is simply difficult to predict for a given UE what SINR changes will occur when a given interferer is silenced. The changes in SINR attributed to silencing a given interferer are not uniform across the cell and the degree of improvement in SINR is in fact geographically specific.

To examine this geographic specificity, a cell at the centre of a two tier hexagonal deployment was populated with UEs at a resolution of one UE per $7m^2$ and the SINR of each UE was measured under two conditions (a) both immediate neighbours transmitting (b) neither immediate neighbour transmitting. Figure 6-1 plots the difference $b - a$ in SINR observed under these two conditions, to show how the gain in SINR attributed to neighbour silencing, varies across the cell.

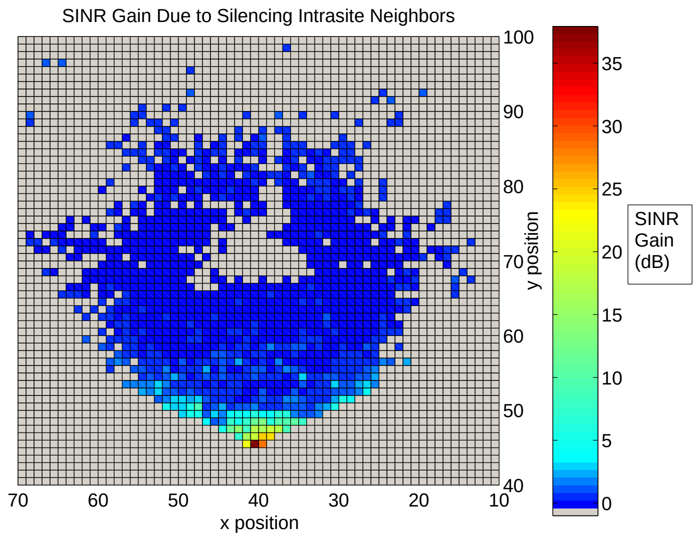

Figure 6-1: Gain in SINR attributed to silencing intrasite neighbours, as a function of geographical position.

In the figure, the antenna boresight has its origin at coordinates $(40, 46)$ and points north. The intrasite neighbours are thus incident upon the bottommost edges which extend to the left and right of the boresight point. The figure clearly shows that significant gains in SINR are seen predominantly along these incident edges, and this is not surprising, for it is at these points that the transmissions from the adjacent intrasite cells will have the most impact. The CDF associated with this cell is shown in Figure 6-2.

About 5% of UE positions benefit from improvements in SINR of 10dB or greater. It is interesting to ask whether or not these gains are substantial enough to overcome the bandwidth loss associated with enforcing the intrasite orthogonality required to obtain them.

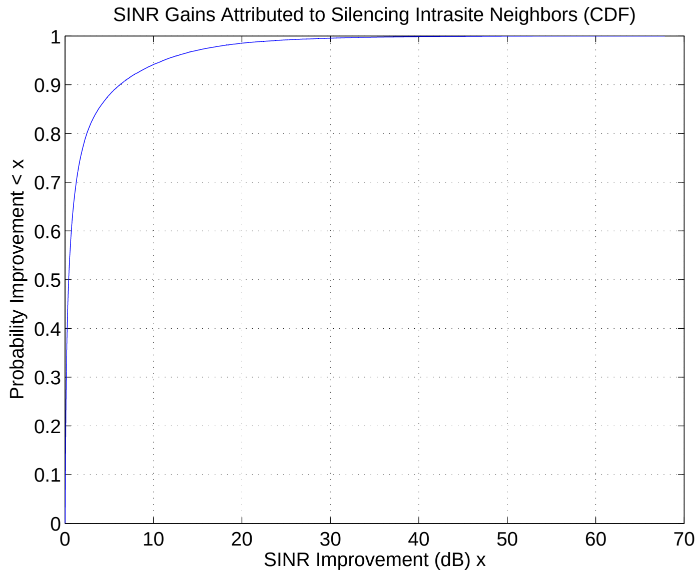

Figure 6-2: CDF of gains attributed to silencing intrasite neighbours.

Figure 6-3 shows the mapping between SINR and BPS (Bits Per Symbol) generated by the Vodafone R&D link level simulations. The solid line shows a realistic estimation of the mapping between SINR and BPS, assuming the availability of 2x2 MIMO. The dotted line is identical to the solid line but scaled by a factor of 3. It therefore shows how much of an increase in BPS would be required to offset a 1/3rd reduction in BW. This can then be mapped to how much of an increase in SINR would be required by tracing the horizontal line.

By performing this tracing automatically, for each input SINR, the minimum SINR improvement required to double or triple the throughput can be computed. This is plotted in Figure 6-4.

To take an example, for an input SINR of $-6dB$, doubling the bitrate would require an SINR improvement of $\approx 7dB$, whereas tripling the bitrate would require an SINR improvement of $\approx 8dB$. The graph also shows the SINR improvements required for factor 2.1 and 3.1 bitrate improvements, because it is desired to not only equal the bitrate prior to enforcing orthogonality, but to improve upon it.

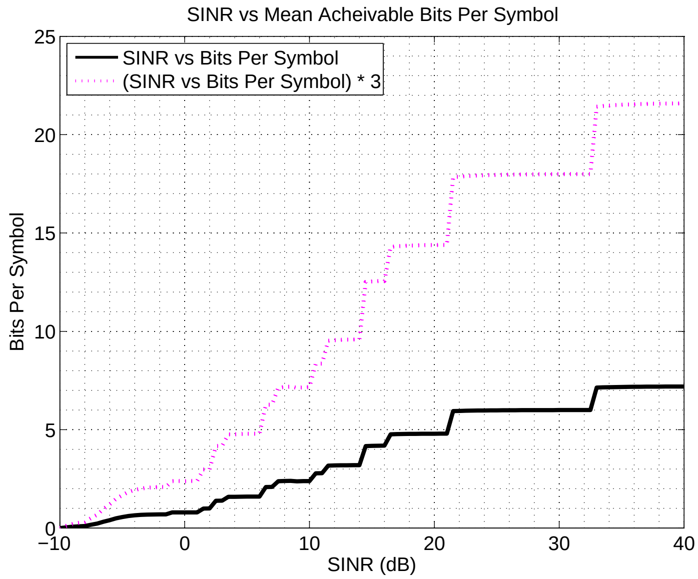

Figure 6-3: Mapping from SINR to BPS and three times this.

Coupled with an ability to predict what actual SINR improvement a given UE is likely to get from enforced orthogonality, the information in Figure 6-4 can be used to decide whether doing so will be beneficial.

If this reasoning is applied to the cell shown Figure 6-1, the BPS gain associated with each geographical position can be plotted. This is shown in Figure 6-5.

As can be seen, the majority of large gains occur near the transmitter, along its incident edges. The percentage of positions with greater than a factor 3 improvement, is small. And from observing the CDF, shown in Figure 6-6, less than 3% of the positions would benefit from mutual intrasite orthogonality.

If the UEs at these positions could be readily identified, and the SINR gains accurately predicted, then enforcing orthogonality should result in an improvement in throughput for these UEs.

Note of course that this preliminary analysis is based on an instantaneous snapshot of only one cell. The actual percentage of UEs that can gain from orthogonality is likely to vary with changes in fast fading, and for different sites. Thus the first step in continuing this work, would be to more accurately quantify the potential of the technique, by averaging

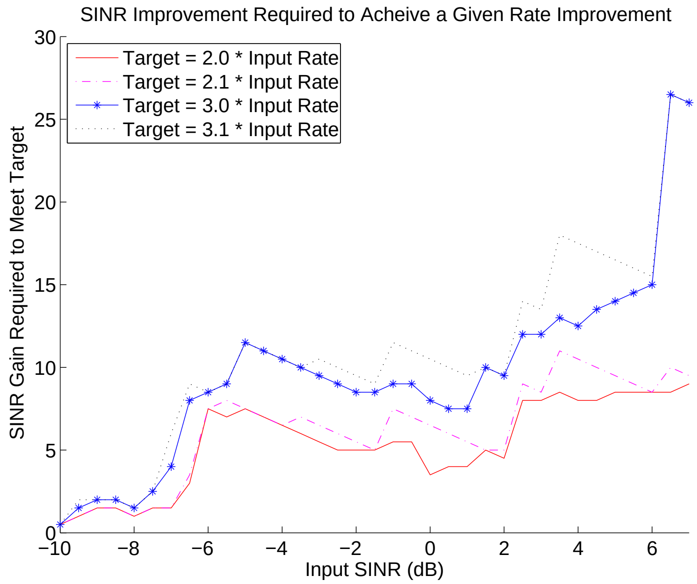

Figure 6-4: SINR improvement required to achieve a given bitrate improvement for various input SINRs.

over multiple TTIs and over multiple shadowing seeds.

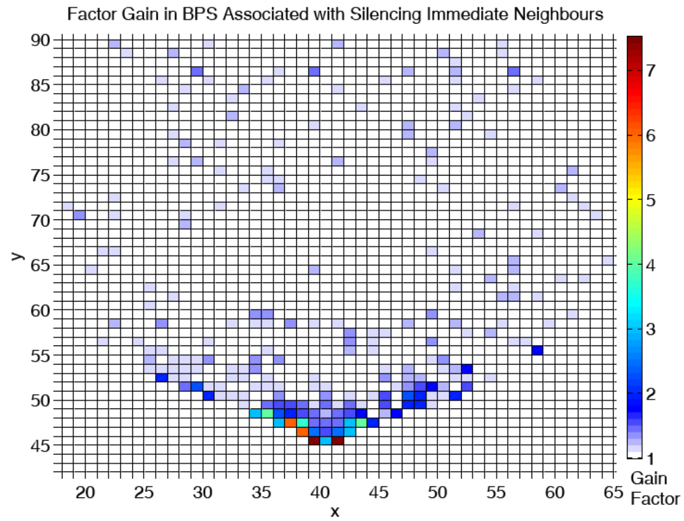

Figure 6-5: BPS Gain associated with different geographical positions when the immediate neighbours are silenced.

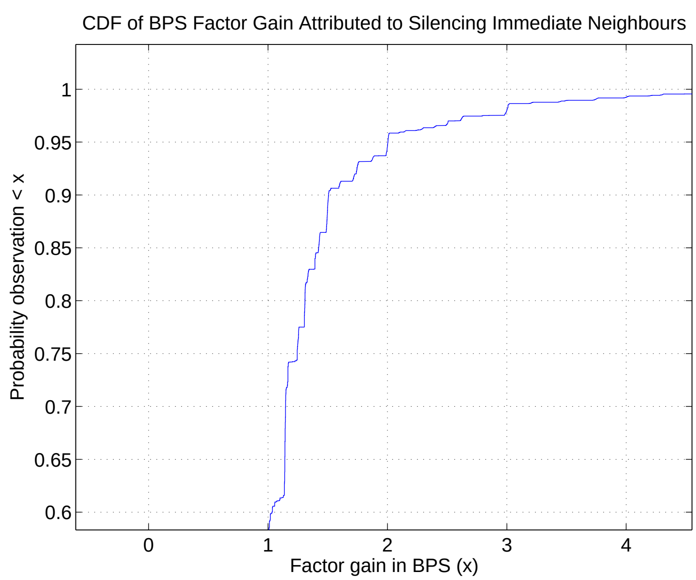

Figure 6-6: CDF of BPS gain, expressed as a factor relative to no neighbour silencing, attributed to silencing the immediate neighbours.

6.2 Predicting CQI

It is worth pointing out that *instantaneous* SINR under full interference is not a reliable predictor of SINR gain under orthogonality. This is because instantaneous SINR varies as much along the intrasite borders as it does anywhere else, and so shows little correlation. A more reliable predictor of SINR gain for a given UE is the RSRPs of the serving sector and the intrasite neighbours coupled with the mean SINR of the serving sector. This contextual information, can be used to predict the likely SINR gain, due to silencing neighbour silencing, for a given UE.

To obtain some confidence that this is true, a feed forward network was trained to predict the SINR gain, due to silencing intrasite neighbours, for each UE measurement obtained in the last section. The structure of the neural network used is shown in Figure 6-7.

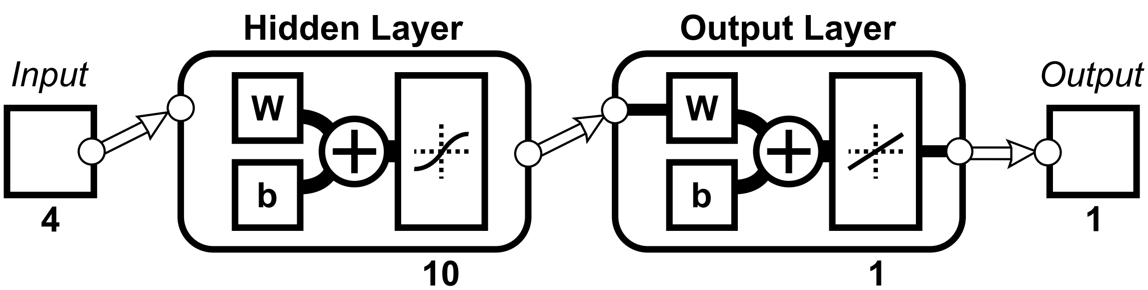

Figure 6-7: Neural network structure. **w** represents weights, and **b** represents the bias weight which in each case has input value 1. The network consists of 4 inputs, 10 sigmoid hidden units, and 1 linear output unit.

The four network inputs are: the SINR of the serving cell, and the RSRPs of the serving cell and the two intrasite neighbours. These inputs are connected to the hidden layer along with a bias. The hidden layer contains 10 units with symmetric sigmoid transfer functions. The hidden layer is connected along with a bias to a single output unit with a linear transfer function. The output produced is the predicted gain associated with silencing the neighbours for the given UE.

The dataset consisted of 74860 instances, with 50% used for training, 25% used for validation, and 25% used for testing. The network weights were trained using Levenberg-Marquardt algorithm [124] with $\mu = 0.001$, a μ decrease ratio of 0.1, a μ increase ratio of 10, and a maximum μ of 1×10^{10}. The algorithm was executed for 1000 epochs using random starting weights. The training performance is shown in Figure 6-8.

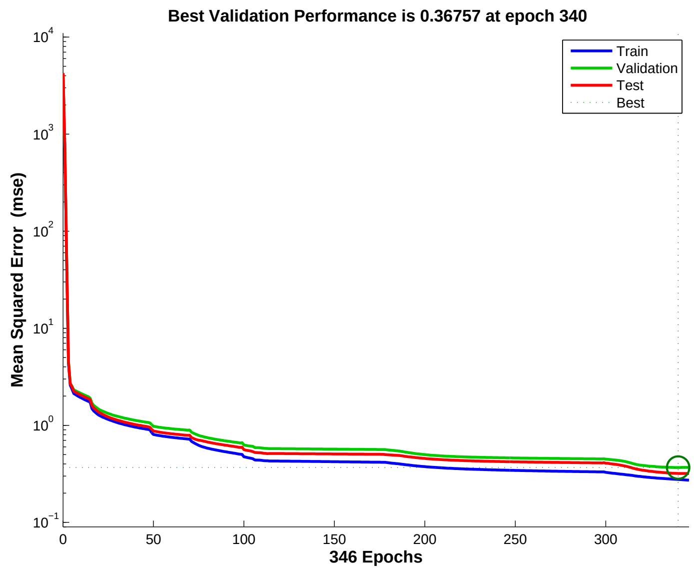

Figure 6-8: Mean Squared Error as a function of training epoch.

The best validation error had an MSE (Mean Squared Error) of 0.36757, which implies a mean prediction error less than $1dB$. Figure 6-9 shows the correlations between the neural network predictions and the training, test, and validation datasets.

As can be seen, the correlations are very good, especially for values $< 30dB$ (the majority of points). From this it can be seen that intrasite RSRPs and serving sector SINR *are* predictive of the gains that occur due to neighbour silencing. Thus correcting CQI based on this information should be possible in practice.

Further work is needed to confirm this, since only a single cell was examined here. In addition, the predictions here are based on only one serving cell SINR measurement for each UE, because the simulation was executed for only one timestep (owing to there being 74860 UEs). In practice of course, the RSRPs and mean SINR will change over time. The results here demonstrate for a given snapshot, that a high prediction accuracy can be obtained, but further work is needed to examine the prediction of SINR gains in dynamic conditions.

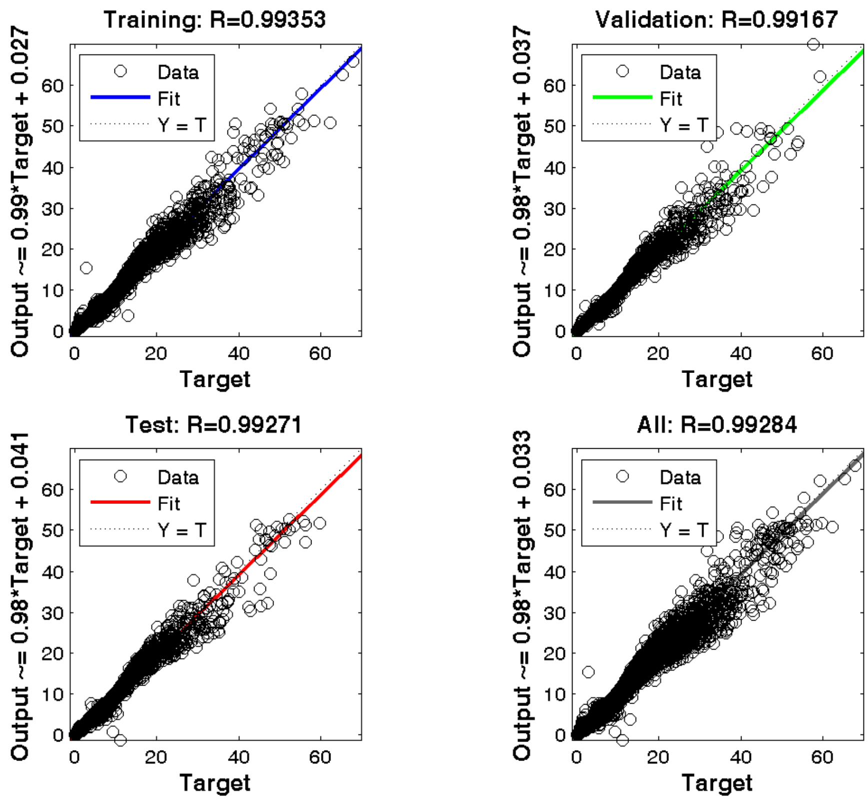

Figure 6-9: Correlations between neural network predictions and objective targets for training, testing, and validation datasets.

6.3 Chapter Summary

This chapter examined how the intrasite scheduling scheme presented in Chapter 5 could be extended to obtain benefits in high interference situations. The concept promoted was to enforce orthogonality between UEs when doing so will result in a net gain. To this end, the improvement due to enforced orthogonality was examined geographically for the cell area, and it was observed that net gains can be obtained for a small percentage of users along the intrasite boundaries, close to the serving sector. To capitalise on this gain a means of predicting the improvement in SINR attributed to the approach was identified as being necessary. To demonstrate that prediction of this improvement is indeed possible, a neural network was trained to perform this prediction, and was shown to obtain low mean squared errors in doing so.

Chapter 7

Conclusions

This chapter begins in Section 7.1 by evaluating the impact of this work in relation to the goals set out in the introduction. Section 7.2 reflects on the processes and conclusions of each of the three reports and discusses how they are related. A summary of the insights provided by this work is provided in Section 7.3, and Section 7.4 goes on to discuss how the quality of work in this domain could be improved by the adoption of some guidelines for best practice. Section 7.5 brings the dissertation to an end by speculating on the likely future of the industry with respect to the work performed here.

7.1 Impact

The introduction of this dissertation set out goals of "homogenising user experience in the future network, in the face of expected diverse user geography, radio conditions, and demand". It is reasonable to ask to what extent these goals have been met.

This dissertation contributes to the understanding of frequency selective scheduling and static interference coordination in Chapters 3 and 4 through extensive simulation and subsequent analysis. This information empowers others with the knowledge to develop novel algorithms, and the analytical approach taken illustrates the need to quantify and identify the source of gains.

As a working example of this, the knowledge gained in Chapters 3 and 4, has been used in Chapter 5 to develop a novel intra-site scheduling scheme that mitigates interference through intelligent coordination of resource usage. The approach performs well in low to moderate load scenarios, and is robust to network conditions within this operating range. It can feasibly be deployed in extant equipment with only moderate vendor cooperation and requires no modification to the standards.

In the area of heavy traffic loads, Chapter 6 paves the way for future research, by quantifying coordination gains as a function of geographical position. Using this information, coupled with a demonstration that it is possible to predict CQI changes under coordination using simple feed forward networks, the stage is set for exploitation and further development of intra-site scheduling algorithms in heavy load conditions.

In this regard, this dissertation has identified and developed practical means of improving the homogeneity of user experience at low to moderate loads, and identified the means by which improvements in the same can be achieved at high loads.

7.2 Reflection

The next section reflects on the scientific processes employed by this dissertation. Section 7.2.2, which follows, comments on the extent to which the outcomes of the work in Chapter 3, Chapter 4, and Chapter 5 can be unified.

7.2.1 Process

This dissertation has chosen experimental procedures matched to the level of abstraction appropriate to the questions posed. In Chapter 3, the focus was on evaluating decoder performance for different subcarrier selection mechanisms at various UE velocities. In order to prevent the misinterpretation of chance aberration as a true result, statistical averaging was performed over many decoder blocks for each control point examined. The results presented however did not *quantify* the degree of confidence within which the observed differences could be expected to lie.

Some explanation is required as to why this was so. The business need was to produce data to help understand the behaviour of WiMAX in a timeframe consistent with the ongoing standards activities. The generation of the link level curves is notoriously time intensive, and this work is no exception. The link level curves generated for the work in Chapter 3 required the simulated transmission and decoding of $6,960,000$ code blocks. This took approximately one month of computer time on 40 XEON 5130 cores each running at 2.0Ghz. For reasonable statistical tests of significance to be performed, it is estimated that a further six months of computation would have been required. It was therefore decided to forego statistical rigour in favour of interpreting the results against the computer time actually expended, and to bound the scope of interpretation accordingly in line with the expected value of the information outcome.

The value of the information produced in Chapter 3 was in being able to challenge vendor simulations being presented at the time, and since the questions posed were principally concerned with the performance of decoders, detailed link level simulations were unavoidable. This is in contrast to the work of Chapter 4 and Chapter 5 which principally address system level aspects of cellular networks. In the latter case, further link level simulation was not required because the work in Chapter 3 had already established some confidence in the efficacy of frequency selective scheduling at the targeted UE velocities, in line with results presented by others. The underlying link level results were also found to be compatible with vendor produced curves. In Chapter 4 and 5 link-level curves produced by vendors from testing hardware in the field were therefore adopted, since they were now considered accurate enough.

In Chapter 4 and Chapter 5 the computational effort for the most part far exceeded that of the competing literature. This is because it was computationally practical to produce enough data for sound statistical tests to be performed; the latter being an element which the majority of competing work lacked. The amount of computational effort expended for the work in Chapter 4 was approximately twice that expended in Chapter 3, but many more results were produced. The work in Chapter 5 applied about the same computational effort as in Chapter 4 but the execution time was shortened due to the availability of a greater number of processing cores.

The modelling policy applied by this dissertation can be summarised as follows: Given the expected variance of the observables under study, we estimated the number of data instances needed to compute the likelihood that observed differences were due to chance.

This estimate set the bounds for the *ideal* number of simulation runs and thus the time window within which such a set of runs could be expected to fall. These temporal bounds were weighed against the expected value of the resultant information, in order to decide whether it was worth investing the ideal simulation effort.

In the case of link-level simulations performed for Chapter 3, the "ideal" set of simulations (as defined here) would have taken longer than the useful commercial lifetime of the generated information. In this case, the strategy taken was to: (i) simulate as many blocks as possible, to maximise confidence in the link-level results. (ii) Compare the generated link-level results against public data and data obtained from partners as a means of validation.

In contrast, for the simulations performed in Chapter 4 and Chapter 5, the weighing of the expected commercial value of the information against the expected time required to generate it, came out in the positive direction. In that, it was decided to take the time to generate the ideal set of data. Such that statistical tests for significance could be performed.

Note that making this kind of effort is in the minority when viewed in the context of the wider literature, but is necessary if full scientific integrity is to be maintained. This is not to say however, that the information generated by others in curtailing this effort is useless. But it certainly makes it a lot harder to evaluate how useful it is.

We believe that it is always necessary to *at least* pass comment on the quality of obtained data, and the extent to which this might affect the accuracy the inferences made. This is something that many extant publications in this field fail to do.

In conclusion, the extensive simulation effort exerted by this work was integral to the modelling strategy. A policy of matching computational effort to the expected commercial and academic value of the work was adopted, wherever it was not feasible to generate enough data for statistical tests to be performed. We will now briefly reflect on the commonality of results among the three chapters.

7.2.2 Common Threads

In Chapter 3 the performance of frequency selective scheduling was examined under differing UE velocities. The performance benefits attributed to frequency selective scheduling were shown to degrade with increasing velocity, eventually to the detriment of the approach in favour of the compared frequency spreading approach. Thus, not only does successful exploitation of channel information depend on its immediacy, but, acting on the basis of stale information can in fact be harmful. This is an important point to consider, since the work in Chapter 4 and Chapter 5 was conducted under an assumption of accurate channel quality indications, i.e. low UE velocities.

More generally the message received is that changing network conditions can lead an otherwise aptly chosen scheduling strategy to perform not only badly, but worse than *a previously inferior strategy*. Meaning that the best strategy for a given scenario must be deftly chosen to match the observed conditions. This message is reinforced in Chapter 4, where particular interference avoidance strategies are shown to perform well for given scenarios, badly in others, with none performing generally well across the board.

This observation itself bars many interference strategies from consideration, precisely because they *cannot* adapt to changing conditions. The desire for simple static solutions yields to the dynamism apparent in real networks. The benefits of tracking dynamism are enforced in Chapter 5 where an algorithm is presented which reacts to both local conditions and the conditions of neighbouring cells.

Being aware of the accuracy of information is critical. And we see that the averaging strategy (frequency spreading) in Chapter 3, which prevailed against frequency selective scheduling at high velocities, was designed explicitly to handle this lack of accurate channel information, with the goal that no single user should suffer more than any other on average. It is this knowledge of the lack of knowledge, which leads to a better result in that case. The conclusion therefore, must be that excellent network performance depends not only on receiving state information, but also in understanding where its limitations lie. In this sense, each of the reports are tied together in their illustration of this point.

7.3 Insights

Intelligent resource scheduling is one of the most important factors in determining cell throughput and throughput fairness. Scheduling subsumes the majority of gains obtained by static interference mitigation solutions, obviating the need for it, and leaves 5% to 10%

of UE positions amenable to gains through intrasite resource orthogonality.

At low to medium loads, intrasite scheduling can readily exploit these gains by leveraging the additional frequency and user diversity that manifests. At high loads, especially in full interference scenarios, gains are obtainable only through enforced orthogonality or coordinated transmit powers.

Intrasite gains are small in magnitude, relative to the opportunistic scheduling gain already obtained over the application of static schemes. The gains from coordination with non-intrasite neighbours are likely to apply to a higher percentage of UE positions, and are likely to be larger that those available from intrasite coordination alone. This follows from the observation that SINR will improve whenever an additional interferer is identified and mitigated, and increasing the area of consideration will increase the number of potential interferes that can be considered. The uncertain factors are: whether or not these interferers can be identified, and whether or not the mitigation gains will exceed the equivalent costs of the additional coordination effort required.

Future work should focus on simple dynamic schemes that enforce orthogonality between selected UEs, in different cells whenever this can reliably identified as having a benefit. The extent of this benefit should be quantified and published in analytical work, through sampling of realistic UE distributions.

Work is needed to improve CQI estimation methods to address inaccuracies that arise from the on-off transmit nature of other cells. This should be integrated into a prediction framework to reliably identify the SINR improvements attributed to any given UE, when any given interferer is silenced, and used to produce reliable interference mitigation solutions.

Future algorithms should be as simple as possible, so as to increase the chance of standardisation or the inclusion of standardised features to support them. Work should be backed up with analysis that facilitates understanding and generalisation of the presented techniques.

With the advent of LTE Advanced, greater opportunity will be met with greater complexity; researchers should look forward to solutions that mitigate the latter and exploit the former.

7.4 Best practices

In Section 2.3 (Page 78), many simulation works were observed that did not consider frequency selective scheduling as a benchmark when championing their novel schemes. This made the results less useful than they could be.

It would excellent if the industry had a set of guidelines for experiments, or a list of recommended simulation assumptions. And too some extent it does, a 3GPP document exists [70] which outlines a number of recommended simulation assumptions including: antenna model, cell layout, signal propagation, details of the eNB and those of the UE, example AWGN based link level curves, and MCS sets. It also describes methodology for simulation of coexistence scenarios.

It falls short, in not expanding upon this to describe methodologies for other scenarios, such as interference coordination, handover, scheduling, beamforming, CoMP, and so on. There is no calibration scenario provided, which could be used to align simulators against a common ground, in order to make results more comparable. A significant number of useful recommendations are thus emitted from the document, such that researchers typically invent their own scenarios, and fill in the gaps using their own expert knowledge. In this sense, it is difficult to reconcile the multitude of results seen in the literature.

There are a number of reasons why clear guidelines do not exist. First of all, because so many different research questions are posed, it is difficult to judge how to create a finite set of scenarios that satisfies everyone's needs, if this is even possible.

Secondly, every research group typically has there own simulation platform, and there is a wide variety of detail and abstraction seen. Whilst it is desirable to create an open platform, and there are a number of projects with this goal in mind (See for example openWNS [125]) it is again difficult to create a platform that satisfies everyone's needs. This does not of course prohibit the idea of creating multiple simulation platforms to alleviate this problem, but how many would be needed? And how would results between them be reconciled?

Another major area in which the support literature falls short, is in the availability of empirical measurements of real equipment and networks. For example, it is difficult to come by accurate traffic data, that shows how traffic changes as a function of time, and how traffic patterns and artifacts such as hotspots emerge and decay. This is partly due to operators being reluctant to disclose certain information to their competitors. An example which applies more directly to the work in this dissertation, is the lack of measurements for intrasite correlation. It would be great if industry and academia could collaborate to create a realistic and open data sets for use by researchers.

Finally, a greater emphasis needs to be placed on analysis and exposition, rather than simple ranking of one technique in comparison to another. This has already been discussed for the case of *trade-offs* in interference mitigation, so another example will be provided.

In [111] and [99] interference mitigation techniques are proposed which place restrictions on resources at a temporal granularity larger than that at which scheduling decisions are made. It is not clear why restrictions are placed on resources based on their impact to other cells, when that impact changes at a rate faster than that at which the restrictions are placed. There seems to be an assumption that fast fading is at some temporal scales, fractally self similar. But there is not any justification or evidence provided to support this. And even if *the model* displays such self-similarity, this could be an artifact that is not seen in reality.

7.5 The Future of the Industry

Caveat: there is a saying usually attributed to Mark Twain and professed notably by Neils Bohr, that "Prediction is difficult, especially of the future". Whilst this is a joke, one need look no further than the continued existence of the stock market, to find incontrovertible evidence that the future cannot be predicted accurately by all of us, all of the time.

The mobile industry is at present in a state of transition between 3G services and the next generation services. LTE Release 8 is being deployed, with roll-out occurring first in urban areas, followed by on-demand rural roll-out.

Within LTE Release 8, vendors are likely to promote proprietary interference avoidance packages and schedulers, and sell these on the basis of indicated results. It is probable that this will be done without clear explanation of the underlying approaches, in order to protect intellectual property. The best way that researchers can get their ideas into Release 8, is to file patents that cover the likely solutions before others do, and collaborate with startup companies that are offering RRM services where possible.

Looking beyond LTE Release 8, the network is set to undergo significant changes. LTE Advanced [126] (LTE Release 10) is the natural candidate for network evolution and is likely to be adopted by Vodafone. LTE Advanced intends to include a number of features that will have a direct bearing on interference and its mitigation. Some examples [127, 128] are described below:

- Beam forming and higher order MIMO

  Beam forming is a manner in which multiple antennas can be exploited. In beamforming the antennas are phased to direct the boresight to the transmission target, or to direct a null at some interferer. In [95] it was shown that beamforming can improve throughput significantly. Beamforming can be considered an interference mitigation scheme in itself, but is one a number of ways that multiple antennas can be used. The other major way is to use the antennas for MIMO, and in LTE Advanced it is proposed that up to 8 layer transmission be adopted in the downlink and 4 in the uplink.

- CoMP (Coordinated Multi Point) transmission and reception

  CoMP techniques leverage the X2 interface to improve throughput by exchanging information about transmissions. For example: a number of eNB can receive the UL transmission of some UE and share that information with the serving eNB, thereby improving detection. In another method, called Successive Interference Cancellation (SIC), cells exchange information about their uplink transmissions such that interference can be “subtracted” upon reception [128]. More generally, CoMP techniques can exploit multi-point transmission, joint processing, coordinated scheduling, and coordinated beamforming. Techniques are likely to appear first for the intrasite case due to simplicity.

- Relays

  In relaying, a transmission from source to target goes via one or more intermediate transmitters. This is promoted as being able to improve coverage, reduce power consumption, and is useful for areas where wired backhaul is unavailable. If relaying becomes a prominent feature in future networks, interference coordination is likely to appear in two manners: coordination within a given cell between eNB and relay, and joint consideration of relay use among neighbouring cells. For example, it is easy to imagine there might be an optimal choice for two relays for a given pair of UEs, that trades off the mutual interference with the improved signal.

- Heterogeneous (multilayer) networks

  Future networks are liable to contain a wide variety of different cell sizes and sources of interference. This makes the interference environment more complicated but also increases the degrees of freedom with which to combat it.

- Carrier aggregation

  In traditional deployments, for example a UMTS FDD deployment, it is typical to use paired spectrum such that one spectrum allocation is used for downlink, and the other for uplink. Carrier aggregation allows more than one spectrum allocation to be used for downlink and uplink transmissions, taken from a pool of either contiguous or disjoint bands. This is relevant to interference coordination, because having a choice of carriers provides more opportunities for exploitative coordination.

Many of these additions affect interference in a similar manner: they increase the degrees of freedom within which interference mitigation can operate. This is a double-edged sword since on the one hand there are likely to be more opportunities for gain, but on the other hand, the greater complexity will make it more difficult to exploit these opportunities.

Another major area in which the future of the industry is set to change, is in energy consumption. The UK Government has agreed to cut carbon emissions relative to a 1990 baseline by 80% by 2050 and by at least 34% by 2020 [129]. Vodafone has its own target to reduce CO2 emissions by 50% by 2020, relative to the 2007 financial year, in all operating companies obligated under the Kyoto protocol [130].

Vodafone R&D is actively engaged in this process, and is an industrial partner in the Mobile VCE [131] Green Radio program, which aims to obtain a "100-fold reduction in power consumption over current wireless communication networks" whilst minimising the impact on quality of service and cost [132].

The additional constraint of being sensitive to energy demands, is likely to increase the complexity of interference mitigation. Interestingly however, it in some sense the goals are complementary: under reduced interference, lower transmit powers can be tolerated for a given MCS target. Conversely, under lower transmit power, greater resilience to interference is required.

Bibliography

[1] Press Release: The future belongs to the mobile internet, Accessed 08/03/2011. http://www.vodafone.com/content/index/press/local_press_releases/germany/germany_press_release/the_future_belongs.html.

[2] International Council on Systems Engineering: A Consensus of the INCOSE Fellows: Commentary, Accessed 03/02/2011. http://www.incose.org/practice/fellowsconsensus.aspx.

[3] Alexander Kossiakoff and William N. Sweet. *Systems Engineering Principles and Practice (Wiley Series in Systems Engineering and Management)*. Wiley-Blackwell, 2002. ISBN 978-0470405482.

[4] Robert Shishko. *NASA Systems Engineering Handbook, Revision 1, NASA/SP-2007-6105*. NASA, 2005. ISBN 978-1780391380.

[5] Peter Checkland. *Systems Thinking, Systems Practice*. John Wiley & Sons, 1981. ISBN 978-0471986065.

[6] Peter Checkland and Jim Scholes. *Soft Systems Methodology In Action*. John Wiley & Sons, 1990. ISBN 978-0471986058.

[7] 3GPP TS 45.001 - V9.0 - Radio Access Network; Physical layer on the radio path; General description. http://www.3gpp.org/ftp/Specs/archive/45_series/45.001/45001-900.zip, May 2009.

[8] Siegmund M. Redl and Matthias K. Weber and Malcolm W. Oliphant. *An introduction to GSM*. Artech House Publishers, 1995. ISBN 978-0890067857.

[9] Timo Halonen, Jaview Romero, and Juan Melero. *GSM, GPRS, and EDGE Performance*. Wiley, 2002. ISBN 978-0470866948.

[10] 3GPP TS 25.308 - V10.1.0 - High Speed Downlink Packet Access (HSDPA); Overall description; Stage 2 (Release 10). http://www.3gpp.org/ftp/Specs/archive/25_series/25.308/25308-a10.zip, 2010.

[11] Andrew J. Viterbi. *CDMA: principles of spread spectrum communication*. Addison-Wesley, 1995. ISBN 978-0201633740.

[12] Harri Holma and Antti Toskala. *WCDMA for UMTS: HSPA Evolution and LTE*. Wiley-Blackwell, 5th edition, 2010. ISBN 978-0470686461.

[13] Third Generation Partnership Project (3GPP) Homepage, Accessed 03/02/2011. http://www.3GPP.org.

[14] Ofcom Homepage, Accessed 03/02/2011. http://www.ofcom.org.uk.

[15] The Wireless Telegraphy (Third Generation Licences) Notice 1999. http://www.ofcom.org.uk/static/archive/spectrumauctions/InformationMemorandum/Downloads/Files/RevNotice.pdf.

[16] John S. Oakland. *Total Quality Management: The route to improving performance*. Butterworth Heinemann, 2nd edition, 1993. ISBN 978-0750609937.

[17] Yoji Akao. *Quality Function Deployment: Integrating Customer Requirements Into Product Design*. Productivity Press, 2004. ISBN 978-1563273131.

[18] QFD ONLINE... moving into the House of Quality: Free QFD Templates, Accessed 03/02/2011. http://www.qfdonline.com/templates.

[19] QFD Tutorial. http://elsmar.com/pdf_files/QFD-Tutorial.swf.

[20] Paul Sandle, Reuters Group Limited. O2 says iPhone demand strained its London network, Accessed 03/02/2011. http://www.reuters.com/article/idCNLDE5BS0M520091229, December 2009.

[21] MONOTAS: Mobile Network Optimisation Through Advanced Simulation Homepage, Accessed 03/02/2011. http://www.macltd.com/monotas.

[22] Ashley Mills, David Lister, Marina De Vos and Yusheng Ji. The Impact of MS velocity on The Performance of Frequency Selective Scheduling in IEEE 802.16e Mobile WiMAX. In *Proceedings of the 7th Annual IEEE Consumer Communications & Networking Conference (IEEE CCNC 2010)*. Las Vegas, Nevada USA, January 2010.

[23] SOCRATES Project Homepage, Accessed 03/02/2011. `http://www.fp7-socrates.eu/`.

[24] Neil Scully and Stefan Thiel and Remco Litjens and Ljupco Jorguseski and Renato Nascimento and Ove Linnell and Kristina Zetterberg and Mehdi Amirijoo and Chris Blondia and Kathleen Spaey and Ingrid Moerman and Irina Balan and Thomas Krner and Andreas Hecker and Thomas Jansen and Jakub Oszmianski and Lars Christoph Schmelz. INFSO-ICT-216284 SOCRATES D2.1: Use Cases for Self-Organising Networks. `http://www.fp7-socrates.eu/files/Deliverables/SOCRATES_D2.1Usecasesforself-organisingnetworks.pdf`.

[25] Ashley Mills and David Lister. Intrasite Scheduling for Interference Avoidance in LTE. In *To appear in proceedings of: 2011 IEEE 73rd Vehicular Technology Conference: VTC2011-Spring, 15-18 May 2011, Budapest, Hungary*, 2011.

[26] M.P.M. Hall and L.W. Barclay and M.T. Hewitt, editor. *Propagation of Radiowaves*. Institution of Electrical Engineers, 1996. ISBN 978-0852961025.

[27] V.S. Abhayawardhana and I.J. Wassell and D. Crosby and M.P. Sellars and M.G. Brown. Comparison of Empirical Propagation Path Loss Models for Fixed Wireless Access Systems. In *Vehicular Technology Conference, 2005. VTC 2005-Spring. 2005 IEEE 61st*, volume 1, pages 73–77, June 2005.

[28] European Conference of Postal and Telecommunications Administrations. SEAMCAT Handbook. Technical report, European Communications Office, January 2010.

[29] Guidance on selecting the propagation model, Accessed 03/02/2011. `http://tractool.seamcat.org/wiki/Manual/PropagationModels/Introduction`.

[30] Bernard Sklar. Rayleigh Fading Channels in Mobile Digital Communication Systems Part I: Characterization. *IEEE Communications Magazine*, 35(7):90–100, 1997.

[31] Danian S. Baum, Jan Hansen, Giovanni Del Galdo, Marko Milojevic, Jari Salo, and Pekka Kyoesti. An Interim Channel Model for Beyond-3G Systems: Extending the 3GPP Spatial Channel Model (SCM). In *IEEE 61st Vehicular Technology Conference*, volume 5, pages 3132–3136. Stockholm, Sweden, May 2005.

[32] Henrik Schulze and Christian Lüders. *Theory and applications of OFDM and CDMA: Wideband Wireless Communications*. John Wiley & Sons, 2005. ISBN 978-0470850695.

[33] Mischa Schwartz. *Mobile Wireless Communications*. Cambridge University Press, 2005. ISBN 978-0521843478.

[34] Volker Kühn. *Wireless Communications over MIMO Channels: Applications to CDMA and Multiple Antenna Systems*. John Wiley & Sons, 2006. ISBN 978-0470027165.

[35] S.M. Alamouti. A simple transmit diversity technique for wireless communications. *Selected Areas in Communications, IEEE Journal on*, 16(8):1451 –1458, Oct 1998.

[36] Markus Rupp and Christoph F. Mecklenbrauker. On Extended Alamouti Schemes for SpaceTime Coding. In *In Proceedings of 5th International Symposium on Wireless Personal Multimedia Communications (WPMC '2002), October 27-30, Hawaii, USA*, 2002.

[37] Loutfi Nuaymi. *WiMAX: Technology for Broadband Wireless Access*. John Wiley & Sons Ltd, 2007. ISBN 978-0470028087.

[38] David J.C. MacKay. *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, 2003. ISBN 978-0521642989.

[39] Pokhariyal, A. and Kolding, T.E. and Mogensen, P.E. Performance of Downlink Frequency Domain Packet Scheduling for the UTRAN Long Term Evolution. In *Personal, Indoor and Mobile Radio Communications, 2006 IEEE 17th International Symposium on*, September 2006.

[40] Pokhariyal, A. and Pedersen, K.I. and Monghal, G. and Kovacs, I.Z. and Rosa, C. and Kolding, T.E. and Mogensen, P.E. HARQ Aware Frequency Domain Packet Scheduler with Different Degrees of Fairness for the UTRAN Long Term Evolution.

In *Vehicular Technology Conference, 2007. VTC2007-Spring. IEEE 65th*, pages 2761–2765, April 2007.

[41] Pokhariyal, A. and Monghal, G. and Pedersen, K.I. and Mogensen, P.E. and Kovacs, I.Z. and Rosa, C. and Kolding, T.E. Frequency Domain Packet Scheduling Under Fractional Load for the UTRAN LTE Downlink. In *Vehicular Technology Conference, 2007. VTC2007-Spring. IEEE 65th*, pages 699–703, April 2007.

[42] Mongha, G. and Pedersen, K.I. and Kovacs, I.Z. and Mogensen, P.E. QoS Oriented Time and Frequency Domain Packet Schedulers for The UTRAN Long Term Evolution. In *Vehicular Technology Conference, 2008. VTC Spring 2008. IEEE*, pages 2532–2536, May 2008.

[43] Kian Chung Beh and Simon Armour and Angela Doufexi. Joint Time-Frequency Domain Proportional Fair Scheduler with HARQ for 3GPP LTE Systems. In *Vehicular Technology Conference, 2008. VTC 2008-Fall. IEEE 68th*, pages 1 –5, Sep. 2008.

[44] Nararat Ruangchaijatupon and Ji Yusheng. Simple Proportional Fairness Scheduling for OFDMA Frame-Based Wireless Systems. In *Wireless Communications and Networking Conference, 2008. WCNC 2008. IEEE*, April 2008.

[45] Assaad, M. and Mourad, A. New Frequency-Time Scheduling Algorithms for 3GPP/LTE-like OFDMA Air Interface in the Downlink. In *Vehicular Technology Conference, 2008. VTC Spring 2008. IEEE*, pages 1964–1969, May 2008.

[46] Ashish Pandharipande and Marios Kountouris and Ho Yang and Hyoungwoon Park. Subcarrier Allocation Schemes for Multiuser OFDM Systems. In *SPCOM 2004, 7th Biennial International Conference on Signal Processing and Communications*. Samsung Advanced Institute of Technology Communications and Networking Lab, December 2004.

[47] Ericsson. The Evolution of EDGE, White Paper. `www.ericsson.com/res/docs/whitepapers/evolution_to_edge.pdf`, February 2007.

[48] International Telecommunication Union (ITU) Homepage, Accessed 03/02/2011. `http://www.itu.int`.

[49] International Mobile Telecommunications-2000 (IMT-2000) Homepage, Accessed 03/02/2011. `http://www.itu.int/home/imt.html`.

[50] Third Generation Partnership Project 2 (3GPP2) Homepage, Accessed 03/02/2011. http://www.3GPP2.org.

[51] European Telecommunications Standards Institute (ETSI) Homepage, Accessed 03/02/2011. http://www.etsi.org.

[52] Association of Radio Industries and Businesses (ARIB) Homepage, Accessed 03/02/2011. http://www.arib.or.jp/english.

[53] Telecommunication Technology Committee (TTC) Homepage, Accessed 03/02/2011. http://www.ttc.or.jp/e.

[54] Alliance for Telecommunications Industry Solutions (ATIS) Homepage, Accessed 03/02/2011. http://www.atis.org.

[55] Telecommunications Technology Association (TTA) Homepage, Accessed 03/02/2011. http://www.tta.or.kr/English.

[56] China Communications Standards Association (CCSA) Homepage, Accessed 03/02/2011. http://www.ccsa.org.cn/english.

[57] Telecommunications Industry Association (TIA) Homepage, Accessed 03/02/2011. http://www.tiaonline.org/.

[58] IEEE. The IEEE 802.16 Working Group on Broadband Wireless Access Standards Homepage, Accessed 03/02/2011. http://ieee802.org/16.

[59] Wimax forum homepage, accessed 03/02/2011. http://www.wimaxforum.org.

[60] Intel Corporation Homepage, Accessed 03/02/2011. http://www.intel.com.

[61] Clearwire Homepage, Accessed 03/02/2011. http://www.clearwire.com.

[62] Sprint Homepage, Accessed 03/02/2011. http://www.sprint.com.

[63] 3GPP TS 36.300 - V8.5.0 - Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Overall description; Stage 2 (Release 8). http://www.3gpp.org/ftp/Specs/archive/36_series/36.300/36300-850.zip, May 2008.

[64] P802.16Rev2/D4 April 2008, DRAFT Standard for Local and metropolitan area networks Part 16: Air Interface for Broadband Wireless Access Systems. Consolidation of IEEE 802.16-2004, 802.16e-2005, 802.16f-2005, and 802.16g-2007.

[65] Krishna Balachandran and Doru Calin and Fang-Chen Cheng and Niranjan Joshi and Joseph H. Kang and Achilles Kogiantis and Kurt Rausch and Askok Rudrapatna and James P. Seymour and Jonqyin Sun. Design and analysis of an IEEE 802.16e-based OFDMA communication system. *Bell Labs technical journal*, 11:53–73, 2007.

[66] WiMAX Forum. WiMAX Forum Mobile System Profile Release 1.0 Approved Specification (Revision 1.4.0: 2007-05-02).

[67] Seung Joon Lee. Trade-Off Between Frequency Diversity Gain and Frequency-Selective Scheduling Gain in OFDMA Systems with Spatial Diversity. *Communications Letters, IEEE*, 11(6):507–509, 2007.

[68] N. Riato, S. Sorrentino, D. Franco, C. Masseroni, M. Rastelli, and R. Trivisonno. Impact of Mobility on Physical and MAC Layer Algorithms Performance in Wimax System. *IEEE 18th International Symposium on Personal, Indoor and Mobile Radio Communications. PIMRC 2007.*, pages 1–6, September 2007.

[69] Nikola Tesla. Method of Signaling: Specification forming part of Letters Patent No. 723,188, dated March 17, 1903.

[70] 3GPP TS 36.942 - V8.2.0 - Evolved Universal Terrestrial Radio Access (E-UTRA); Radio Frequency (RF) system scenarios. `http://www.3gpp.org/ftp/Specs/archive/36_series/36.942/36942-820.zip`, May 2009.

[71] Andras Racz and Norbert Reider and Gabor Fodor. On the Impact of Inter-Cell Interference in LTE. In *Global Telecommunications Conference, 2008. IEEE GLOBECOM 2008. IEEE*, pages 1–6, 2008.

[72] A Simonsson. Frequency Reuse and Intercell Interference Co-Ordination In E-UTRA. In *Vehicular Technology Conference, 2007. VTC2007-Spring. IEEE 65th*, pages 3091 –3095, April 2007.

[73] Thomas Toftegard Nielsen and Jeroen Wigard and Preben Mogensen. On the capacity of a GSM frequency hopping network with intelligent underlay-overlay. In

Vehicular Technology Conference, 1997 IEEE 47th, volume 3, pages 1867–1871, May 1997.

[74] Siemens. 3GPP TSG RAN WG1 Meeting #41, R1-050476: Evolved UTRA uplink scheduling and frequency reuse. `http://www.3gpp.org/ftp/tsg_ran/wg1_rl1/TSGR1_41/Docs/R1-050476.zip`, 2005.

[75] Huawei. 3GPP TSG RAN WG1 Meeting #41, R1-050507 - Soft Frequency Reuse Scheme for UTRAN LTE. `http://www.3gpp.org/ftp/tsg_ran/wg1_rl1/TSGR1_41/Docs/R1-050507.zip`, 2005.

[76] Ericsson. 3GPP TSG RAN WG1 Meeting #42, R1-050764 - Inter-cell Interference Handling for E-UTRA. `http://www.3gpp.org/ftp/tsg_ran/wg1_rl1/TSGR1_42/Docs/R1-050764.zip`, 2005.

[77] LG Electronics. 3GPP TSG RAN WG1 Meeting #42, R1-050833 - Interference mitigation in evolved UTRA/UTRAN. `http://www.3gpp.org/ftp/tsg_ran/wg1_rl1/TSGR1_42/Docs/R1-050833.zip`, 2005.

[78] Alcatel. 3GPP TSG RAN WG1 Adhoc Meeting, R1-050593 - Interference coordination for evolved UTRA uplink access. `http://www.3gpp.org/ftp/tsg_ran/wg1_rl1/TSGR1_AH/LTE_AH_June-05/Docs/R1-050593.zip`, 2005.

[79] Texas Instruments. R1-051059: Inter-Cell Interference Mitigation for EUTRA. `http://www.3gpp.org/ftp/tsg_ran/wg1_rl1/TSGR1_42bis/Docs/R1-051059.zip`, October 2005.

[80] Alcatel. 3GPP TSG RAN WG1 #42: R1-050694: Multi-cell Simulation Results for Interference Co-ordination in new OFDM DL. `http://www.3gpp.org/ftp/tsg_ran/wg1_rl1/TSGR1_42/Docs/R1-050694.zip`, September 2005.

[81] Siemens. R1-050738: Interference mitigation - Considerations and Results on Frequency Reuse. `http://www.3gpp.org/ftp/tsg_ran/wg1_rl1/TSGR1_42/Docs/R1-050738.zip`.

[82] Mohammad Abaii et al. IST-4-027756. WINNER II D4.7.2 v1.0. Interference avoidance concepts.

[83] Nokia. 3GPP TSG RAN WG1 Meeting #44, R1-060291: OFDMA Downlink Inter-cell Inteference mitigation. `http://www.3gpp.org/ftp/tsg_ran/wg1_rl1/TSGR1_44/Docs/R1-060291.zip`, 2006.

[84] F. Xiangning, C. Si, and Z. Xiaodong. An Inter-Cell Interference Coordination Technique Based on Users Ratio and Multi-Level Frequency Allocations. In *Wireless Communications, Networking and Mobile Computing (WiCom 2007)*, pages 799–802, September 2007.

[85] Gary Boudreau and John Panicker and Ning Guo and Rui Chang and Neng Wang and Sophie Vrzic. Interference Coordination and Cancellation for 4G Networks. *IEEE Communications Magazine*, 47(4), April 2009.

[86] Shuping Chen, Wenbo Wang, Xing Zhang, and Wei Hong. Performance and Optimization of Interference Coordination in Downlink TDD-LTE Systems. In *Information Management and Engineering (ICIME), 2010 The 2nd IEEE International Conference on*, pages 181–185, April 2010.

[87] Xiang Yikang, Luo Jijun, and Christian Hartmann. Inter-cell Interference Mitigation through Flexible Resource Reuse in OFDMA based Communication Networks. In *In proceedings of 13th European Wireless Conference (EW'07), Paris, France*, April 2007.

[88] Siemens. 3GPP RAN WG1 Ad Hoc Meeting, R1-060135 - Interference Mitigation by Partial Frequency Reuse. `http://www.3gpp.org/ftp/tsg_ran/wg1_rl1/TSGR1_AH/LTE_AH_January-06/Docs/R1-060135.zip`, 2006.

[89] IPWireless. 3GPP RAN WG1 Ad Hoc Meeting, R1-050652 - Attaining the Cell Edge Performance Requirements for the LTE Downlink. `http://www.3gpp.org/ftp/tsg_ran/wg1_rl1/TSGR1_AH/LTE_AH_June-05/Docs/R1-050652.zip`, 2005.

[90] Texas Instruments. 3GPP TSG RAN WG1 Meeting #44, R1-060368 - Performance of Inter-Cell Inteference Mitigation with Semi-Static Frequency Planning for EUTRA Downlink.txt. `http://www.3gpp.org/ftp/tsg_ran/wg1_rl1/TSGR1_44/Docs/R1-060368.zip`, 2006.

[91] Ericsson. R1-061374: Downlink inter-cell interference co-ordination/avoidance evaluation of frequency reuse. http://www.3gpp.org/ftp/tsg_ran/wg1_rl1/TSGR1_45/Docs/R1-061374.zip, May 2006.

[92] Jonathan A C Sterne and George Davey Smith. Sifting the evidencewhat's wrong with significance tests? *BMJ*, 322(7280), January 2001.

[93] NTT DoCoMo Ericsson. TSG-RAN WG1 meeting #44, R1-060586 - Downlink and uplink inter-cell inteference co-ordination/avoidance - impact on the specifications. http://www.3gpp.org/ftp/tsg_ran/wg1_rl1/TSGR1_44/Docs/R1-060586.zip, 2006.

[94] R.Y. Chang, Zhifeng Tao, Jinyun Zhang, and C.-C.J. Kuo. Multicell ofdma downlink resource allocation using a graphic framework. *Vehicular Technology, IEEE Transactions on*, 58(7):3494 –3507, sep. 2009.

[95] M.C. Necker. Towards Frequency Reuse 1 Cellular FDM/TDM Systems. In *Proceedings of the 9th ACM/IEEE International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM 2006)*, October 2006.

[96] Marc C. Necker. ocal Interference Coordination in Cellular OFDMA Networks. In *Vehicular Technology Conference, 2007. VTC-2007 Fall. 2007 IEEE 66th*, pages 1741–1746, October 2007.

[97] Marc C. Necker. Integrated scheduling and interference coordination in cellular ofdma networks. In *Broadband Communications, Networks and Systems, 2007. BROADNETS 2007. Fourth International Conference on*, pages 559 –566, sep. 2007.

[98] Marc C. Necker. Scheduling Constraints and Interference Graph Properties for Graph-based Interference Coordination in Cellular OFDMA Networks. *Mobile Networks and Applications*, 14(4):539–550, 2009.

[99] Guoqing Li and Hui Liu. Downlink Radio Resource Allocation for Multi-Cell OFDMA System. *Wireless Communications, IEEE Transactions on*, 5(12):3451–3459, December 2006.

[100] Gabor Fodor and Chrysostomos Koutsimanis and Andrs Rcz and Norbert Reider and Arne Simonsson and Walter Muller. Intercell Interference Coordination in OFDMA

Networks and in the 3GPP Long Term Evolution System. *Journal of Communications*, 4(7):445–453, August 2009.

[101] Mahmudur Rahman and Halim Yanikomeroglu. Interference Avoidance through Dynamic Downlink OFDMA Subchannel Allocation using Intercell Coordination. In *Vehicular Technology Conference, 2008. VTC Spring 2008. IEEE*, pages 1630–1635, May 2008.

[102] Christos Papadimitriou and Kenneth Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Dover Publications, 1998. ISBN 978-0486402581.

[103] Mahmudur Rahman and Halim Yanikomeroglu and William Wong. Interference Avoidance with Dynamic Inter-Cell Coordination for Downlnink LTE System. In *WCNC'09: Proceedings of the 2009 IEEE conference on Wireless Communications & Networking Conference*, pages 1238–1243. IEEE Press, Piscataway, NJ, USA, 2009. ISBN 978-1-4244-2947-9.

[104] Rahman, M. and Yanikomeroglu, H. Enhancing cell-edge performance: a downlink dynamic interference avoidance scheme with inter-cell coordination. *Wireless Communications, IEEE Transactions on*, 9(4):1414–1425, April 2010.

[105] 3GPP TS 36.423 V8.9.0 - Evolved Universal Terrestrial Radio Access Network (E-UTRAN); X2 application protocol (X2AP) (Release 8). `http://www.3gpp.org/ftp/Specs/archive/36_series/36.423/36423-890.zip`, March 2010.

[106] 3GPP TS 36.213 V8.8.0 - Evolved Universal Terrestrial Radio Access (E-UTRA); Physical layer procedures (Release 8). `http://www.3gpp.org/ftp/Specs/archive/36_series/36.213/36213-880.zip`, September 2009.

[107] Xuehong Mao and Amine Maaref and Koon Hoo Teo. Adaptive Soft Frequency Reuse for Inter-Cell Interference Coordination in SC-FDMA Based 3GPP LTE Uplinks. In *Global Telecommunications Conference, 2008. IEEE GLOBECOM 2008. IEEE*, pages 1 –6, December 2008.

[108] A. Persson, T. Ottosson, and G. Auer. Inter-Sector Scheduling in Multi-User OFDM. In *IEEE International Conference on Communications, 2006. ICC '06.*, volume 10, pages 4415–4419, 2006.

[109] Kun Dong and Hui Tian and Xingmin Li and Qiaoyun Sun. A Distributed Inter-Cell Interference Coordination Scheme in Downlink Multicell OFDMA Systems. In *Consumer Communications and Networking Conference (CCNC), 2010 7th IEEE*, pages 1–5, January 2010.

[110] He Xiaoben and Valkealahti, K. On Distributed and Self-Organized Inter-Cell Interference Mitigation for OFDMA Downlink in IMT-Advanced Radio Systems. In *Vehicular Technology Conference, 2007. VTC-2007 Fall. 2007 IEEE 66th*, pages 1736–1740, October 2007.

[111] Shuqin Zheng and Hui Tian and Zheng Hu and Lan Chen and Jianchi Zhu. QoS-Guaranteed Radio Resource Allocation with Distributed Inter-Cell Interference Coordination for Multi-Cell OFDMA Systems. In *Vehicular Technology Conference (VTC 2010-Spring), 2010 IEEE 71st*, pages 1–5, May 2010.

[112] Soumaya Hamouda and Sami Tabbane and Philippe Godlewski. Improved reuse partitioning and power control for downlink multi-cell OFDMA systems. In *ACM International Conference Proceeding Series; Vol. 196*, volume 196. ACM Press, Alghero, Italy, September 2006.

[113] Alcatel-Lucent. R1-073187: Interference Coordination Framework with Results, June 2007.

[114] J. Ellenbeck, C. Hartmann, and L. Berlemann. Decentralized inter-cell interference coordination autonomous spectral reuse decisions. In *14th European Wireless Conference (EW 2008)*, pages 1–7, June 2008.

[115] Alexander L. Stolyar and Harish Viswanathan. Self-organizing Dynamic Fractional Frequency Reuse for Best-Effort Traffic Through Distributed Inter-cell Coordination. In *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, pages 691–699, April 2008.

[116] Alexander L. Stolyar and Harish Viswanathan. Self-organizing Dynamic Fractional Frequency Reuse Through Distributed Inter-cell Coordination: The Case of Best-Effort Traffic. Technical report, Bell Labs, Alcatel-Lucent, Technical Memo, May 2008.

[117] Francis Dominique et al. Self-Organizing Interference Management for LTE. *Bell Labs Technical Journal*, 15(3):19–42, 2010.

[118] Nageen Himayat and Shilpa Talwar and Anil Rao and Robert Soni. Interference management for 4G cellular standards [WIMAX/LTE UPDATE]. *Communications Magazine, IEEE*, 48(8):86 –92, August 2010.

[119] Nokia and Nokia Siemens Networks. 3GPP TSG RAN WG1 #49, R1-072974, Downlink interference coordination. `http://www.3gpp.org/ftp/tsg_ran/wg1_rl1/TSGR1_49b/Docs/R1-072974.zip`, 2007.

[120] R Bosisio and U Spagnolini. Interference Coordination Vs. Interference Randomization in Multicell 3GPP LTE System. In *Wireless Communications and Networking Conference, 2008. WCNC 2008. IEEE*, pages 824–829, 2008.

[121] Daniele Molteni and Monica Nicoli. Interference Mitigation in Multicell LTE Systems: Performance over Correlated Fading Channels. In *Spread Spectrum Techniques and Applications, 2008. ISSSTA '08. IEEE 10th International Symposium on*, pages 133 –139, August 2008.

[122] Sternad, M. and Ottosson, T. and Ahlen, A. and Svensson, A. Attaining both coverage and high spectral efficiency with adaptive OFDM downlinks. In *Vehicular Technology Conference, 2003. VTC 2003-Fall. 2003 IEEE 58th*, volume 4, pages 2486–2490, October 2003.

[123] S. Kumar, G. Monghal, J. Nin, I. Ordas, K. I. Pedersen, and P. E. Mogensen. Autonomous Inter Cell Interference Avoidance under Fractional Load for Downlink Long Term Evolution. In *69th Vehicular Technology Conference (VTC Spring 2009), Barcelona, Spain*, pages 26–29, April 2009.

[124] Christoper M. Bishop. *Neural Networks for Pattern Recognition.* Claredon Press, 1995. ISBN 978-0198538646.

[125] Home - openWNS (open Wireless Network Simulator). Accessed 03/02/2011. `http://www.openwns.org`.

[126] Takehiro Nakamura, 3GPP TSG-RAN Chairman. Proposal for Candidate Radio Interface Technologies for IMT-Advanced Based on LTE Release 10 and Beyond (LTE-Advanced). `http://www.3gpp.org/IMG/pdf/2009_10_3gpp_IMT.pdf`.

[127] Ghosh, A. and Ratasuk, R. and Mondal, B. and Mangalvedhe, N. and Thomas, T. LTE-advanced: next-generation wireless broadband technology [Invited Paper]. *Wireless Communications, IEEE*, 17(3):10–22, June 2010.

[128] Stencel, V. and Muller, A. and Frank, P. LTE Advanced: A further evolutionary step for Next Generation Mobile Networks. In *Radioelektronika (RADIOELEKTRONIKA), 2010 20th International Conference*, pages 1–5, April 2010.

[129] Climate Change Act 2008 - Department of Energy and Climate Change website, Accessed 03/02/2011. `http://www.decc.gov.uk/en/content/cms/legislation/cc_act_08/cc_act_08.aspx`.

[130] Vodafone's response to stakeholders - Energy use and climate change, Accessed 03/02/2011. `http://www.vodafone.com/content/index/about/sustainability/approach/stakeholder_engagement/vodafone_s_responsetostakeholders.html`.

[131] Mobile VCE Homepage, Accessed 03/02/2011. `http://www.mobilevce.com`.

[132] Virtual Centre of Excellence in Mobile and Personal Communications - Mobile VCE - CORE-5 Research Area: Green Radio. `http://www.mobilevce.com/infosheets/GreenRadio.pdf`.