Norwegian University of
Science and Technology

The Wireless Tram

Providing Wireless Trondheim/Wi-Fi coverage using mobile WiMAX as backhaul

Rein Sigve Karlsen

Master of Science in Communication Technology
Submission date: July 2008
Supervisor: Yuming Jiang, ITEM
Co-supervisor: Thomas Jelle, Trådløse Trondheim

Norwegian University of Science and Technology
Department of Telematics

Problem Description

This assignment should include both a theoretical and a practical part.

The theoretical part should investigate the IEEE 802.16e-2005 amendment and seek to gain thorough understanding of its technical foundation, features, and performance. A comparison with IEEE 802.11 is also to be conducted.

The practical part should implement a proof-of-concept wireless backhaul using mobile WiMAX. The WiMAX network should feed a Wi-Fi hotspot attached to a moving vehicle, such as a tram, thus providing Internet access to travelling passengers. Tests are to be carried out to investigate the quality of such a service, including bandwidth and signal quality measurements.

Assignment given: 15. January 2008
Supervisor: Yuming Jiang, ITEM

Preface

This thesis was written as the final part of a master's degree in communication technology with specialisation in access and core networks. It was carried out at the Department of Telematics (ITEM) at the Norwegian University of Science and Technology (NTNU) during spring 2008. The thesis covers 30 credits.

In mid 2007 I was introduced to Thomas Jelle, managing director of Wireless Trondheim. He leads a start-up company that seeks to provide Wi-Fi access to the inner city of Trondheim. Besides providing Internet access to its users, the wireless network also acts as a test-bed for new applications and services. During autumn 2007, I carried out a project that focused on the quality of service in wireless LANs with emphasis on testing the new IEEE 802.11e amendment. While working on the project, I talked with Thomas about the possibilities of using WiMAX as an enabler for providing Internet access on public transportation such as buses and trams. We agreed that it would be interesting to gain more knowledge in the field, especially since the new standard for mobile WiMAX had been finalised and products were about to become available.

In addition to Wireless Trondheim, UNINETT also had interests in the new technology. UNINETT develops and operates the Norwegian research network and supplies network services to universities, university colleges, and research institutions. Both companies saw WiMAX as an enabler for providing wireless access outside the reach of Wi-Fi, and were eager to gain more knowledge in the field.

Acknowledgements

Many people have contributed and supported me during the work of this thesis.

First of all, I would like to thank my co-supervisor, Thomas Jelle, managing director of Wireless Trondheim, for the valuable discussions and helpful guidelines that he provided throughput the project. Thanks for all the time you spent with me and for giving me the possibility to undertake such an interesting and extensive project.

My main supervisor, Professor Yuming Jiang, also deserves thanks for valuable support, time, and guidance.

Many thanks also go to a number of persons who have contributed in various ways. Aslaug Hagestad Nag, Bjarte Sørhus, and the other folks at The Norwegian Post- and Telecommunication Authority (PT) for their time and willingness to letting me carry out radio planning using their facilities. Geir Øien at the Department of Electronics and Telecommunications (IET) for valuable feedback related to MIMO, diversity, and signal propagation. Tekna and Atle Tangedal for letting me join the WiMAX conference in Oslo and present my project. Roald Boge and Svorka Aksess for introducing me to their work with testing WiMAX in motion. The people at NextNet for providing valuable information regarding their WiMAX network and experience in the field. Gråkallbanen for permission to perform testing on their trams while in service. Endless hours were spent on their comfy trams, an experience not easily forgotten. Stefano Bertelli and Terje Haugen at the Department of Engineering Cybernetics (ITK) for use of their car while performing measurements in the field.

And last but not least, friends and fellow students for good times, technical discussions, and long hours of Zatacka - off working hours naturally. The Ræga office still rules big time.

Abstract

In spring 2008, a pre-mobile WiMAX network is installed in Trondheim. Its purpose is to cover the tram line and provide wireless backhaul to a Wi-Fi hotspot inside the tram, thus providing Internet access to travelling passengers. Testing is carried out in order to investigate if WiMAX can be used in such scenario. Based on radio propagation modelling, two sectors are initially set up to cover the entire tram line. The WiMAX system operates in the 2.6 GHz band with a bandwidth of 5 MHz.

Testing reveals that reception along the first half of the tram line, covered by the first sector, is mostly according to the Okumura-Hata models for open and sub-urban areas. Adequate signal quality for a high grade of service (-80 dBm) is achieved up to 1.5 km from the base station. The second half of the tram line, covered by the second sector, is mostly according to the Okumura-Hata models for sub-urban and urban areas. The coverage provided by the second site is only enough to provide adequate signal quality half a kilometre from the base station. It is concluded that two sectors are not adequate to cover the entire tram line from start to end.

The impact of 2^{nd} order diversity is also investigated with respect to signal quality and throughput. Separation in space and by polarisation are both tested. A noticeable increase in average throughput of 50% is measured when using diversity with separation in space. Separation by polarisation yields an average increase of 14% only. Improvements in coverage when using diversity is very small, measured to a few hundred meters.

It is noted that the throughput achieved along the tram line varies with the speed of the tram. When standing still, maximum TCP throughput in the downlink is measured to just above 6 Mbps, which is only 70% of the theoretical maximum of 9 Mbps when using an UL/DL ratio of 50/50. While in motion, the TCP throughput seldom rises above 1 Mbps. It is expected that the low throughput in motion is caused by a combination of Doppler effects and the use of an omni-directional antenna on the tram.

Contents

List of Figures

List of Tables

List of Abbreviations

1PPS	one pulse per second
3G	third generation
3GPP	Third Generation Partnership Project
3GPP2	Third Generation Partnership Project 2
AAS	adaptive antenna systems
AC	alternating current
ACK	acknowledgement
ADSL	asymmetric digital subscriber line
AES	advanced encryption standard
AMC	adaptive modulation and coding
ARQ	automatic repeat-query
ASN	access service network
ASN-GW	ASN gateway
ASP	application service provider
ATM	asynchronous transfer mode
ATPC	automatic transmission power control
AU-IDU	indoor access unit
AU-ODU	outdoor access unit
AVU	air ventilation unit
BDP	bandwidth-delay product
BE	best-effort
BER	bit error rate
BPSK	binary phase-shift keying
BS	base station
BTC	block turbo codes
CBC-MAC	cipher block chaining message authentication code
CC	convolutional coding
CC	chase combining
CCI	co-channel interference
CCM	counter mode with CBC-MAC

CDD	cyclic delay diversity
CDMA	code division multiple access
CID	connection identifier
CMAC	cipher-based message authentication code
CP	cyclic prefix
CPE	customer premises equipment
CQICH	channel quality information channel
CRC	cyclic redundancy check
CS	convergence sub-layer
CSMA	carrier sense multiple access
CSN	connectivity service network
CTC	convolutional turbo codes
DC	direct current
DECT	digital enhanced cordless telecommunications
DFT	discrete Fourier transform
DL	downlink
DOA	direction of arrival
DOCSIS	data over cable service interface specification
DP	decision point
DSL	digital subscriber line
DVB-T	digital video broadcasting - terrestrial
EAP	extensible authentication protocol
EIRP	equivalent isotropically radiated power
EP	enforcement point
ertPS	extended real-time polling service
ETSI	European Telecommunications Standards Institute
EV-DO	evolution data optimised
FBSS	fast base station switching
FCH	frame control header
FDD	frequency division duplexing
FDM	frequency division multiplexing
FDMA	frequency division multiple access
FEC	forward error correction
FFT	fast Fourier transform
FSPL	free space path loss
FUSC	full usage of sub-carriers
Gbps	gigabits per second
GMH	generic MAC header
GPS	global positioning system
GSM	global system for mobile communications
HARQ	hybrid ARQ

H-FDD	half-duplex FDD
HHO	hard handover
HIPERMAN	high performance radio metropolitan area network
HMAC	hash-based message authentication code
HSDPA	high-speed downlink packet access
HSUPA	high-speed uplink packet access
ICI	inter-carrier interference
IEEE	Institute of Electrical and Electronics Engineers
IET	Department of Electronics and Telecommunications
IETF	Internet Engineering Task Force
IFFT	inverse fast Fourier transform
IP	internet protocol
IR	incremental redundancy
ISI	inter-symbol interference
ISM	industrial, scientific and medical
ISP	Internet service provider
ITEM	Department of Telematics
ITK	Department of Engineering Cybernetics
ITU	International Telecommunication Union
Kbps	kilobits per second
LAN	local area network
LDPC	low-density parity check
LLC	logical link control
LMDS	local multipoint distribution system
LOS	line-of-sight
LTE	long term evolution
LWAPP	lightweight access point protocol
MAC	media access control
Mbps	megabits per second
MBS	multicast and broadcast services
MBWA	mobile broadband wireless access
MDHO	macro diversity handover
MIB	management information base
MIMO	multiple-input multiple-output
MIP	mobile IP
MIR	maximum information rate
MISO	multiple-input single-output
MMDS	multi-channel multi-point distribution services
MPDU	MAC protocol data unit
MPLS	multi-protocol label switching
MRRC	maximum receive ratio combining

MS	mobile station
MSDU	MAC service data unit
NACK	non-acknowledgement
NAP	network access provider
NLOS	non-line-of-sight
NMS	network management system
NPU	network processing unit
NRM	network reference model
nrtPS	non-real-time polling service
NSB	Norwegian State Railways
NSP	network service provider
NTNU	Norwegian University of Science and Technology
NWG	network working group
OFDM	orthogonal frequency-division multiplexing
OFDMA	orthogonal frequency-division multiple access
OHOA	Okumura-Hata for open areas
OHSA	Okumura-Hata for sub-urban areas
OHUA	Okumura-Hata for urban areas
OSI	Open Systems Interconnection
PAPR	peak-to-average power ratio
PHS	packet header suppression
PHY	physical
PIU	power interface unit
PKM	privacy and key management protocol
PSK	phase-shift keying
PSU	power supply unit
PT	The Norwegian Post- and Telecommunication Authority
PUSC	partial usage of sub-carriers
QAM	quadrature amplitude modulation
QoS	quality of service
QPSK	quadrature phase-shift keying
RADIUS	remote authentication dial-in user service
REP-RSP	report response
RF	radio frequency
RSSI	received signal strength indication
RTG	receive/transmit transition gap
rtPS	real-time polling service
RTT	round-trip time
SC	selection combining
SDMA	spatial division multiple access
SFID	service flow identifier

SIM	subscriber identity module
SIMO	single-input multiple-output
SIR	signal-to-interference ratio
SISO	single-input single-output
SM	spatial multiplexing
SNMP	simple network management protocol
SNR	signal-to-noise ratio
SOFDMA	scalable OFDMA
SPWG	service provider working group
SRTM	shuttle radar topography mission
SSID	service set identifier
STBC	space/time block code
SU	subscriber unit
SU-IDU	subscriber indoor unit
SU-ODU	subscriber outdoor unit
TCP	transmission control protocol
TD-SCDMA	time division synchronous CDMA
TDD	time division duplexing
TDM	time division multiplexing
TDMA	time division multiple access
TTA	Telecommunications Technology Association
TTG	transmit/receive transition gap
UDP	user datagram protocol
UGS	unsolicited grant service
UHF	ultra high frequency
UL	uplink
UMTS	universal mobile telephone system
USIM	universal SIM
VHF	very high frequency
VoIP	voice over IP
VV	second order diversity using separation in space
WCDMA	wideband CDMA
Wi-Fi	wireless fidelity
WiBro	wireless broadband
WiMAX	worldwide interoperability for microwave access
WISP	wireless internet service provider
WLAN	wireless LAN
WLC	wireless LAN controller
WLL	wireless local-loop
WRAN	wireless regional area network
X	second order diversity using separation by polarisation

Chapter 1

Introduction

1.1 Motivation

More and more people are getting WLAN-enabled laptops, mobile phones, or other hand-held devices. As they learn to appreciate the benefits of accessing Internet services over such inexpensive wireless connections, they also wish to connect from wherever they go. As Wireless Trondheim provides Wi-Fi coverage to the inner city of Trondheim, access could be provided using their networks. However, there is currently no coverage outside the inner city, nor along bus or tram routes. Extending the Wi-Fi coverage of Wireless Trondheim to include such places would not be realistic, as the complexity and cost of deploying fibre for Wi-Fi backhaul would be prohibitively expensive. Relying on Wi-Fi or other existing technologies for backhaul would have some challenges too.

Instead, a new and promising technology have materialised. Mobile WiMAX already enjoys broad support from leading vendors, and its future look bright with certified equipment soon to become available. It builds on the success of fixed WiMAX, which thus far has been used to provide fixed and nomadic broadband wireless access. With mobility added, users on the move are easily supported. The promises of mobile WiMAX are appealing. Wireless broadband without a direct line of sight to the base station. Cell radii that may span tens of kilometres. Great performance with throughputs of tens of megabits per second. Low cost deployment. A flourishing ecosystem. On the paper it looks very promising indeed. However, theory and practice are often two different worlds.

In order to verify if WiMAX can deliver what it promises, this thesis first seeks to gain thorough understanding of its technical foundation, and particularly investigate the new functionality present in the mobile WiMAX solution. Based on the findings, a complete WiMAX network is to be acquired and installed in the city of Trondheim. Its usefulness in providing wireless backhaul for a Wi-Fi access point on a moving vehicle, in this case a tram, is to be tested and evaluated. If successful, the solution would enable the delivery of Internet services from Wireless Trondheim to travelling passengers.

1.2 Methodology

The purpose of this thesis is to gain understanding of the WiMAX technology and to investigate its performance when used as backhaul for a moving vehicle, such as a tram. The study of the technology is mostly based on information gathered from reliable sources, namely published books and papers. Some information is also obtained from knowledgeable persons who have expertise in the field.

In order to test the performance of the WiMAX network, an experimental research method has been used. In contrast to an empirical research method, where naturally occurring data is aggregated and analysed, the experimental method strictly manipulates each test in order to observe the changes that result.

1.3 Scope

A number of different parameters may be considered in determining the overall performance of a WiMAX network. In order to limit the scope of this thesis, the main focus is on physical performance and throughput. Physical performance parameters included are among others received signal strength indication, signal to noise ratio, and modulation rate. Throughput is a measure of the data rate that may be achieved, commonly estimated using the UDP or TCP protocols. Both have advantages and disadvantages, however in this thesis only TCP connections have been tested.

There are many other parts of a WiMAX solution that may be investigated. Examples of such are security, quality of service, indoor coverage, and inter-media handover. None of these parameters have been tested in this thesis.

Equipment from a selected WiMAX vendor was tested. It is important to emphasise that this equipment was *not* certified for mobile WiMAX, however it implemented much of the functionality as defined by the IEEE 802.16e-2005 air-interface. Testing was carried out both from fixed locations along the tram line, and on the tram in motion. The client premises equipment used was intended for fixed installations. No mobile stations were tested, as they were not available at the time of writing. Note that testing a different set of equipment, using other frequencies or parameters, or testing at other locations may yield different results than the ones obtained in this project.

1.4 Related Work

The intention of this section is to introduce other related work that has been carried out in the field of WiMAX testing. Most public projects to date have evaluated the performance of fixed WiMAX networks, hence little information exists about mobile WiMAX field trails.

1.4.1 Ball State University

The Office of Wireless Research and Mapping at Ball State University conducted a thorough WiMAX field test during 2007 [1]. The equipment used was an Alvarion BreezeMAX 3500 micro base station operating at 3.5 GHz. A single sector was covered with a 120 degrees sector antenna mounted with 0 degrees downtilt. On the client side, both the BreezeMAX Si CPE, a self-installable indoor subscriber unit, and the BreezeMAX PRO CPE, an outdoor subscriber unit, were tested. A 3.5 MHz channel was used, which allowed for a maximum gross throughput of 12 Mbps.

Many different tests were carried out to investigate the performance of the WiMAX network, and almost 100 different locations were visited to measure the signal strength. The outdoor unit achieved usable signal strength at eight kilometres from the base station while the indoor unit only made it up to one and a half kilometre.

One test evaluated the outdoor signal strength during autumn versus winter. For the outdoor and indoor unit, average differences of 4.1 and 9.3 dB were measured respectively. The study concluded that the signal was stronger during winter than during autumn, mostly due to the difference in tree foliage obstructing the path between the base station and the subscriber unit.

Another test compared the estimated results from using radio propagation modelling software with real measurements. The software was found to predict very accurate results that were close to the actual field test signal readings. A standard deviation of 4.14 dB was found between the two methods. The difference in signal strength between a direct path versus a reflected signal was also tested. The conclusion was that in cases where there is no direct line of sight to the base station, it may be better to rely on a reflected signal. The report demonstrated an example where the reflected signal was quite a few dB better than the direct signal.

Further tests included study of environmental effects such as rain and fog, modulation speed tests and testing of VoIP and video conferencing.

1.4.2 Telenor R&I

During 2006 and 2007, several WiMAX tests were carried out in the area of Oslo, Norway [2].

The first test investigated the performance of a fixed WiMAX network with focus on throughput, coverage, and link quality. Alvarion BreezeMAX 3500 was used as base station, operating on a 3.5 MHz channel in the 3.5 GHz band using FDD. The transmit power of the base station was 28 dBm and the antenna gain was 14 dBi. The Alvarion BreezeMAX PRO CPE was used as subscriber unit, transmitting at 20 dBm with an antenna gain of 18 dBi. Only one CPE was used in the network at any one time. A total of 15 locations were visited during the test, most of them in NLOS conditions. The conclusion of the test was that great performance and signal quality were achieved from sites within 2 km from the base station. Downlink RSSI was measured between -55 dBm and -80 dBm, while for longer distances the RSSI fell below -85 dBm and approached -90 dBm at 11 km. The uplink RSSI achieved results comparable to the downlink RSSI closer than 2 km to the base station. For distances beyond 2 km, the uplink RSSI fell more rapidly mostly due to the limited transmission power of the subscriber unit versus the base station. The throughput varied between 1 and 9.6 Mbps, where the maximum measured was only 76% of the theoretical throughput of 12.71 Mbps. This difference was caused by protocol overhead and management traffic.

A second test investigated the performance of a real life fixed WiMAX network with more than 850 subscribers. A total of ten base stations were deployed, each utilising FDD with a 3.5 MHz channel in both uplink and downlink. The subscriber units used fixed outdoor antennas and were set up with automatic transmission power control (ATPC). Physical performance attributes such as RSSI, SNR, modulation rate, and transmission power were gathered from all the subscriber units every few minutes for a couple of months. Based on the received RSSI values, a path loss model was derived and compared to well known loss models such as the free space path loss (FSPL) and the Cost-231 Hata models for urban and sub-urban areas. The conclusion of the second test was that the path loss experienced in this scenario approached the FSPL model because most of the subscribers had line of sight capabilities to the base station. The Cost-231 Hata models were not applicable for this kind of deployment as the models were developed for mobile systems with NLOS conditions in mind.

The third test investigated a pre-mobile WiMAX system manufactured by Alvarion. The impact of 2^{nd} and 4^{th} order diversity was studied as well as sub-channelisation in the uplink. The latter feature was only partially implemented, as only one subscriber unit could transmit per symbol using groups of 4, 8, or 16 sub-channels. The base station used TDD with a 5 MHz channel in the 3.5 GHz band, and a total of three sectors were set up. Each sector implemented 4^{th} order transmit diversity using cyclic delay diversity (CDD) and 4^{th} order receive diversity using maximum receive ratio combining (MRRC). Two dual-slant antennas were used in each sector, thus both space diversity and polarisation diversity were achieved. The subscriber unit used was Alvarion BreezeMAX Si, a self-installable indoor unit. The third test concluded that highest modulation, hence full throughput, was achieved from NLOS locations within 1 km of the base station. With respect to diversity order, the gain in SNR was greater from 1^{st} to 2^{nd} order diversity, than from 2^{nd} to 4^{th}.

1.4.3 Pilot for NSB

Norwegian State Railways (NSB) is the largest passenger railway company in Norway. During 2007, a project was carried out that investigated the possibility of delivering Internet services on trains using wireless technologies. The area between Skien and Lillehammer was used for the pilot, a distance of about 300 km. Companies involved in the project included Telenor, NetCom, ICE and Svorka Aksess.

Svorka Aksess in co-operation with Nera used WiMAX to cover the desired area. Alvarion BreezeMAX was chosen as supplier for both base stations and CPEs. The equipment was certified for fixed and nomadic use, hence without support for hand-over between cells.

One of the main problems that arose was related to the CPE. The CPE was designed to operate in the network from a fixed location and did not support fast hand-overs between cells. When the CPE moved from one cell to the other, it took a lot of time to re-connect to the new base station. Switching times of up to ten minutes were experienced during the pilot.

The project concluded that relatively high throughput was achieved as long as there was a clear line of sight to the base station, experiencing download speeds of up to 4 Mbps. However, as soon as the CPE entered a NLOS environment, the signal degradation was severe and the throughput dropped considerably. Another noticeable observation was that the signal quality achieved through dense forest with heavy foliage was surprisingly good.

1.4.4 Norwegian Army

The Norwegian Army in co-operation with NextGenTel, Nera, and Telenor R&I performed a fixed and nomadic WiMAX test during fall 2007. Alvarion BreezeMAX TDD equipment was used in the 2.3 GHz band, where each sector implemented 2^{nd} order diversity. Two types of subscriber antennas were tested, both a high gain directional panel antenna with 13 dBi gain and an omni-directional antenna with 4 dBi gain.

The study concluded that the terrain between the base station and subscriber unit had a large impact on the throughput achieved. Measurements performed in dense forest showed varying results, while open areas provided promising throughput. Download speeds of 5 Mbps were achieved at distances of up to 38 km from the base station using the directional antenna. The study also concluded that choosing the most suitable antenna for mobile units is challenging, as directional antennas are difficult to utilise in motion.

1.5 Readers Guide

This thesis is divided into 8 chapters with additional appendices.

Chapter 2 provides an extensive background on the technologies upon which WiMAX is built. Special emphasis is put on the extensions that are included in the mobile WiMAX system profile.

Chapter 3 introduces Wireless Trondheim and details a few interesting scenarios for WiMAX usage.

Chapter 4 deals with the process of network planning. Possible locations for base stations are identified, site surveys are carried out, and propagation modelling is performed to find optimal placement of sites and sectors.

Chapter 5 covers the network implementation phase. The WiMAX solution is outlined and the process of installing and configuring the system is described.

Chapter 6 is devoted to testing of the WiMAX network. All tests that were carried out are thoroughly described here.

Chapter 7 contains a general discussion of the outcome of the testing phase with recommendations for improvement.

Chapter 8 contains conclusions and recommendations for future work.

Chapter 2

Background on WiMAX

This chapter deals with the theoretical part of the project. The first section introduces wireless broadband as a concept and outlines its evolution. Following this, the IEEE 802.16 task group and the WiMAX Forum are presented. Next, the mobile WiMAX system profile is thoroughly described. The network architecture of a complete WiMAX solution is also dealt with before spending some time on comparing WiMAX to other wireless broadband technologies.

2.1 Introduction to Wireless Broadband

Spurred by the deregulation of the Telecom industry and rapid growth of the Internet in the 1990s, competitive carriers and new entrants were motivated to deploy wireless solutions to bypass incumbent service providers and their installed base of fixed wire-line access technologies. During the past decade or so, a number of wireless point-to-multipoint access systems have been developed. According to [3], wireless broadband has evolved through four stages, albeit not fully distinct or clearly sequential:

1. **Narrowband wireless local-loop systems:** Voice telephony was the first application for which a wireless alternative was deployed. The system was called wireless local-loop (WLL), and it later evolved into digital enhanced cordless telecommunications (DECT). As the demand for Internet access increased during the 1990s, several small start-up companies focused on providing Internet-access services to differentiate themselves from voice operators. These wireless Internet service

provider (WISP) companies typically deployed systems capable of a few hundred kilobits per second using licence-exempt 900 MHz and 2.4 GHz bands.

2. **First-generation line-of-sight (LOS) broadband systems:** As digital subscriber line (DSL) and cable modem solutions began to be deployed, wireless systems had to evolve to support much higher speeds to be competitive. The local multipoint distribution system (LMDS) was developed for very high frequencies and supported several hundreds of megabits per second. It showed great promise in the late 1990s, but did not succeed. In early 2000, fixed wireless broadband services were deployed in the multi-channel multi-point distribution services (MMDS) band. Initial systems required both an outdoor antenna and a clear LOS transmission path, which proved to be significant impediments. Also, since each cell covered a relatively large area, the capacity of such systems were fairly limited.

3. **Second-generation non-line-of-sight (NLOS) broadband systems:** The wireless systems evolved to overcome LOS issues and to provide more capacity. By using cellular architectures and implementing advanced signal processing techniques, link and system performance were greatly improved. Throughput of a few megabits per second became possible with these systems.

4. **Standard-based wireless broadband system:** Due to the lack of a common standard, the different solutions developed were not interoperable, and they all varied widely in their performance capabilities, frequency spectrum used, and applications supported. Clearly, the need for standardised and interoperable solutions arose, and this task was taken upon by the Institute of Electrical and Electronics Engineers (IEEE).

The next section introduces the IEEE 802.16 task group and summarises the process of standardising the technology for wireless broadband.

2.2 IEEE 802.16

In 1998, the IEEE formed a group called 802.16 to develop a standard for wireless metropolitan area networks. The primary focus was LOS-based point-to-multipoint solutions in the 10 GHz to 66 GHz band, with a main application of delivering high-speed connections to businesses that could not obtain fibre [3]. The task group produced what is known as the original 802.16 standard [4], which was approved in December 2001. It specified a physical layer that used single-carrier modulation techniques, and supported both frequency division duplexing (FDD) and time division duplexing (TDD). Many of the concepts related to the media access control (MAC) layer were adopted from the popular data over cable service interface specification (DOCSIS).

After completing the standard, the task group continued to expand its capabilities to support operations in licensed and unlicensed frequencies in the 2 GHz to 11 GHz range, which would enable NLOS deployments. One amendment, 802.16a [5], was approved in 2003. It included OFDM as part of the physical layer, and also supported OFDMA on the MAC layer. A further amendment, 802.16c [6], implemented system profiles for the 10 GHz to 66 GHz band. During the same period, the task group also worked on aligning the 802.16 standard with HIPERMAN, which was developed by the European Telecommunications Standards Institute (ETSI). This work was concluded in 2004 with the release of a revised standard, IEEE 802.16-2004 [7], which superseded the earlier 802.16, 802.16a, and 802.16c documents. This standard formed the basis for the first WiMAX deployments, and is referred to as fixed WiMAX.

In December 2005, the 802.16e-2005 amendment [8] was approved. It added support for vehicular mobility and implemented a number of enhancements to the 802.16-2004 standard. This amendment formed the basis for the nomadic and mobile WiMAX solutions.

As it turns out, the 802.16 specifications offer a multitude of design options to accommodate the diverse needs of the industry. There are several physical layer options and a number of choices for duplexing, multiplexing, channel bandwidths, and MAC architecture. One could say that 802.16 is a collection of standards, not a single interoperable standard [3]. Clearly, there was a need for the scope of the standards to be reduced to a smaller set of design choices. Consensus had to be established on what options of the standard to implement and test for interoperability. This task was taken upon by the WiMAX Forum.

The IEEE 802.16 task group continues to improve various aspects of the system. The following sub-sections briefly mention the other 802.16 amendments that are either finalised or being worked on at the moment.

2.2.1 Active Amendments

802.16f-2005 [9] defines a management information base (MIB) for the PHY and MAC layer, along with the associated management procedures. It provides object definitions to enable standard-based management of 802.16 devices using for instance the simple network management protocol (SNMP). The amendment describes a management reference model which consists of a network management system (NMS), managed nodes, and a service flow database.

802.16g-2007 [10] creates a standardised interface for the management of 802.16 devices along with procedures and services for efficient handover, quality of service (QoS) management, and radio resource management.

802.16k-2007 [11] is a relatively short document that amends IEEE 802.1D to support bridging of the 802.16 MAC layer.

2.2.2 Amendments under Development

802.16h seeks to improve the co-existence mechanisms for operation in license-exempt frequency bands. Draft version 7 was released at the end of June 2008.

802.16i updates the MIB where the objective is to add mobility support to the 802.16f fixed MIB. This project was terminated in March 2008 at the request of the 802.16 working group. The content of the 802.16i draft was moved into the 802.16Rev2 project.

802.16j enhances the OFDMA physical layer and MAC layer to enable the operation of relay stations in licensed frequency bands. The purpose of this amendment is to enhance coverage, throughput, and system capacity by specifying multi-hop relay capabilities. The relay station will enable quicker deployment of WiMAX networks, and is expected to be less complex and less expensive than a full-blown base station. The relay station may be used in a number of scenarios. For instance, it could be deployed inside buildings or tunnels to provide better indoor or underground coverage, or it could be mounted on vehicles to provide access on trains, ferries, or buses. Additional features include co-operative diversity gain and relay, where the transmission to or reception from a mobile station is greatly enhanced. The amendment is expected to be finalised at the end of 2010.

802.16m specifies an advanced air-interface that provides performance im-

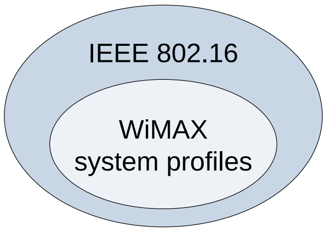

Figure 2.1: WiMAX system profiles

provements necessary to support future services and applications. It seeks to comply with the International Telecommunication Union (ITU) IMT-Advanced requirements for 4G networks, which include 100 Mbps throughput for high mobility, 1 Gbps for low mobility and spectral efficiency of up to 10 bits/sec/Hz. In order to achieve such data rates, a multi-carrier approach is chosen where the transmission takes place on two or more carriers simultaneously. This leads to an increase in bandwidth usage. Also, the 802.16e-2005 frame structure has been shown to be inflexible. Because of this, 802.16m introduces a new frame structure which consists of super-frames which are further divided into frames and again divided into sub-frames. The amendment is expected to be finalised in 2009.

802.16Rev2 consolidates a number of previous standards and amendments into a new document.

The next section introduces the WiMAX Forum and explains its purpose in relation to the 802.16 standards and amendments.

2.3 WiMAX Forum

As mentioned previously, the 802.16 documents contain a plethora of design choices for system developers. For practical reasons of interoperability, there was a need to reduce the scope of these standards and to define a smaller set of design choices for implementation. The WiMAX Forum was thus founded to convert the 802.16 specifications into an interoperable set of system profiles for which products could be certified for conformity and verified for interoperability. In essence, it fulfils the very same purpose as the Wi-Fi Alliance does in relationship to the IEEE 802.11 standards.

The Forum has defined two different system profiles. One is based on the IEEE 802.16-2004 OFDM physical layer and is called the fixed system profile.

Profile name	Frequency band (GHz)	Channel bandwidth (MHz)	Duplexing mode
3.5T1	3.5	7	TDD
3.5T2	3.5	3.5	TDD
3.5F1	3.5	3.5	FDD
3.5F2	3.5	7	FDD
5.8T	3.5	10	TDD

Table 2.1: Fixed WiMAX certification profiles [12]

The other is based on the IEEE 802.16e-2005 scalable OFDMA physical layer and is called the mobile system profile. As illustrated in figure 2.1, a system profile defines the subset of mandatory and optional PHY- and MAC-layer features from the 802.16 standards. Features that are mandatory in the 802.16 documents are also mandatory in the WiMAX system profiles, while optional features may be optional, mandatory, or not included.

Each instantiation of a system profile with a defined operating frequency, channel bandwidth, and duplexing mode is called a certification profile [3]. Currently, there exists five fixed and twelve mobile certification profiles. The different profiles are listed in table 2.1 and 2.2.

2.3.1 Product Certification

The major advantage of standard-based technology is interoperability. However, even though equipment vendors strive to develop products that comply to the standard, specifications may be interpreted and implemented differently [13]. To ensure that equipment from different vendors are indeed interoperable and conform to the standards set forth by the WiMAX Forum, a certification program has been established. The certification process consists of a conformance test and an interoperability test.

The conformance test ensures that the equipment correctly implements all specifications defined by the WiMAX Forum, which is based on the IEEE 802.16 standards. The interoperability test verifies that equipment from different vendors work as expected within the same network. At least three different vendors must submit equipment within the same certification profile to start testing.

The WiMAX Forum operates several certification labs around the globe where vendors may prove that their equipment meet WiMAX standards.

Profile name	Frequency band (GHz)	Channel bandwidth (MHz)	Duplexing mode
MP01	2.3-2.4	8.75	TDD
MP02	2.3-2.4	5	TDD
	2.3-2.4	10	TDD
MP03	2.305-2.320	5	TDD
	2.345-2.360	5	TDD
MP04	2.305-2.320	10	TDD
	2.345-2.360	10	TDD
MP05	2.496-2.690	5	TDD
	2.496-2.690	10	TDD
MP06	3.3-3.4	5	TDD
MP07	3.3-3.4	7	TDD
MP08	3.4-3.8	5	TDD
MP09	3.4-3.6	5	TDD
MP10	3.4-3.6	7	TDD
MP11	3.4-3.8	10	TDD
MP12	3.4-3.6	10	TDD

Table 2.2: Mobile WiMAX certification profiles [12]

Upon successful completion of both tests, the vendors receive a certificate and may market their products as WiMAX Forum Certified.

2.3.2 Certification Releases and Waves

The scope of certification expands through time with the addition of new test cases. In order to ensure backwards compatibility and technology stability, a release and wave framework is used when adding new test cases. Each new release adds new functionality in an incremental fashion, including all the certification tests previously used. This way, backwards compatibility is maintained. As with releases, waves add new test cases within releases, also in an incremental fashion [13].

Fixed WiMAX release 1.0 was approved in 2005 and is divided into two waves. Wave 1 covers mandatory features and includes testing of the air-interface, network entry, dynamic services, and bandwidth allocation. Wave 2 introduces two optional modules, namely QoS and advanced encryption standard (AES). In total, the test suite for fixed WiMAX contains a couple hundred procedures that should be tested.

Mobile WiMAX release 1.0 was approved in February 2006 and is currently divided into two waves as well. Wave 1 covers features such as OFDMA, QoS, AES, key management, support for handover, hybrid ARQ (HARQ), and power control. This wave is ongoing for 2.3 GHz only, all other profiles are starting directly with wave 2 certification. Wave 2 includes advanced features such as multiple-input multiple-output (MIMO), beamforming, and IPv6 in addition to wave 1 features. The depth of testing is greatly increased compared to fixed WiMAX, with a total of 1000 procedures that must be passed. Because of the complexity of the certification process, wave 2 testing has been divided into two phases. As of writing, the WiMAX Forum is currently working on phase 1, while phase 2 is scheduled to commence later in 2008.

The following is a short list of some of the main differences between mobile WiMAX and fixed WiMAX [12, 14]. Many of the techniques and improvements that have been implemented in the mobile WiMAX system profile will be further dealt with in later sections of this thesis.

- Support for mobility with handover procedures
- Power save modes for mobile stations
- Scaling of the fast Fourier transform (FFT) to accommodate different channel bandwidths
- Better implementation of MIMO and adaptive antenna systems (AAS) technology
- Improved channel coding with turbo codes and HARQ
- Sub-channelisation in both uplink (UL) and downlink (DL) to enhance coverage and throughput
- Extra classes for QoS
- Updated security layer with support for a diverse set of encryption and authentication schemes

The rest of this chapter provides more detail about the fixed and mobile WiMAX system profiles. As the focus of this thesis is mobile WiMAX, most emphasis is put on the latter. First, the protocol layers that WiMAX defines are introduced and compared to the OSI reference model. Next, the functions of each layer are dealt with in more detail.

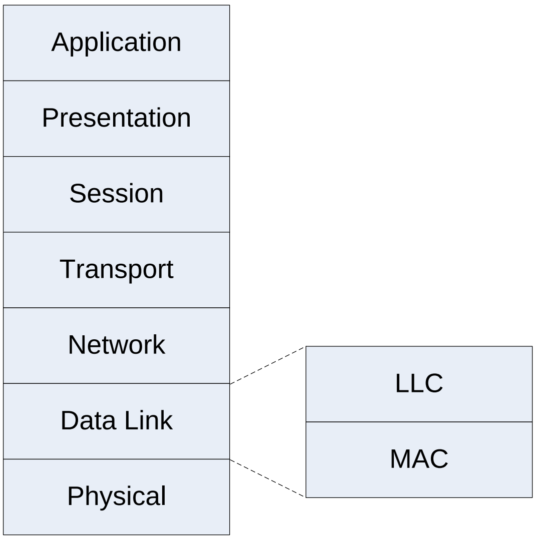

Figure 2.2: The OSI Reference Model [15]

2.4 Protocol Layers

The Open Systems Interconnection (OSI) reference model, as depicted in figure 2.2, is a seven-layered model which is often used to provide an abstract description of a network architecture. Each of the seven layers represents a collection of related functions that provide services to the layer above and receives services from the layer below.

In all networks defined by the IEEE 802 family of standards, the data link layer is split into a logical link control (LLC) sub-layer and a media access control (MAC) sub-layer. The LLC sub-layer is defined by the IEEE 802.2 standard [16] and is common for all IEEE 802 networks. This means that WiMAX, as defined by the IEEE 802.16 standards, specifies only the MAC and physical layers.

2.5 PHY Layer

The IEEE 802.16 suite of standards define five physical layers:

- **WirelessMAN SC**, a single-carrier physical layer intended for frequencies above 11 GHz. It requires LOS conditions and specifies both TDD and FDD duplexing. It is defined in section 8.1 and is part of the original 802.16 standard.

- **WirelessMAN SCa**, a single-carrier physical layer intended for frequencies between 2 GHz and 11 GHz. It is designed for point-to-multipoint operation and specifies both TDD and FDD duplexing. It is defined in section 8.2 of the standard.

- **WirelessMAN OFDM**, a 256-point OFDM physical layer for point-to-multipoint operation intended for NLOS conditions at frequencies between 2 GHz and 11 GHz. It was finalised with the IEEE 802.16-2004 standard and is specified in section 8.3. This physical specification has been accepted by the WiMAX Forum for fixed operations and is referred to as fixed WiMAX.

- **WirelessMAN OFDMA**, a 2048-point OFDMA physical layer for point-to-multipoint operation intended for NLOS conditions at frequencies between 2 GHz and 11 GHz. It is defined in section 8.4 and was completely rewritten in the IEEE 802.16e-2005 amendment to implement scalable OFDMA. This physical specification has been accepted by the WiMAX Forum for mobile and portable operations and is referred to as mobile WiMAX.

- **WirelessHUMAN**, a physical layer for operation in unlicensed frequency bands between 2 GHz and 11 GHz. It is defined in section 8.5 of the standard.

In the following, the functional stages of the WirelessMAN OFDMA physical layer will be explained in more detail. As described earlier, the IEEE 802.16 specifications define a large number of features and options, whereas only a subset of them are included in the WiMAX system profiles. The mobile WiMAX system profile [17], available from the WiMAX Forum, lists which features are mandatory and which are optional.

As illustrated in figure 2.3, the first set of functions are related to the channel coding procedures and include randomisation, forward error correction, bit interleaving, repetition, and modulation. The next set of functions is related to the construction of the OFDM symbol in the frequency domain along with diversity encoding or MIMO if implemented. The final set is responsible for converting the OFDM symbol from the frequency domain to the time domain, and eventually to an analog signal that can be transported over the air. Like all standards, the figure only shows the logical components of the transmitter. Similar components also exist at the receiver to reconstruct the transmitted signal [3].

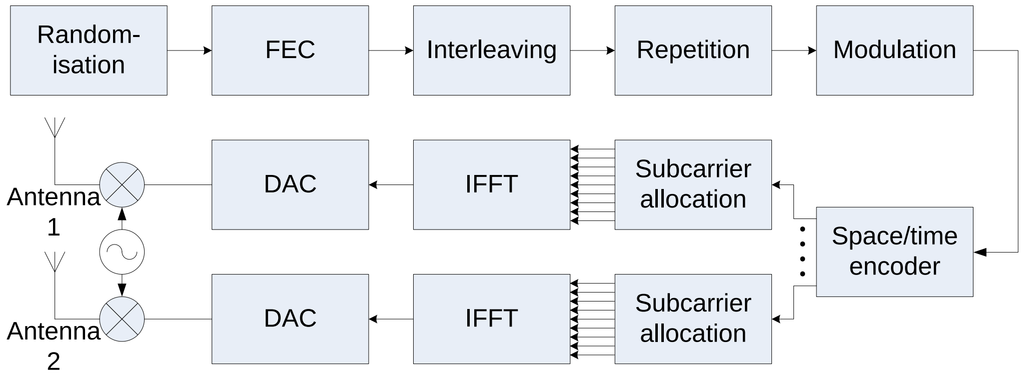

Figure 2.3: Functional stages of the physical layer [7, 8]

2.5.1 Randomisation

Randomisation is performed on all data transmitted on the uplink and downlink, except for the frame control header (FCH). Its purpose is to provide physical layer encryption to prevent a rogue receiver from decoding the data [3]. It also introduces protection in that it avoids long sequences of consecutive zeroes or ones [12]. Randomisation is described in section 8.4.9.1 in the IEEE 802.16 documents and is mandatory for both the base station (BS) and mobile station (MS) in the mobile WiMAX system profile.

2.5.2 FEC

Forward error correction (FEC) is a coding technique whereby the sender adds redundant data to the transmission, allowing the receiver to detect and correct errors. The amount of non-redundant to the total amount of information specifies the code rate, typically given as a fractional number. If the code rate is k/n, then for every k bits of useful information, the encoder generates a total of n bits of data, where $n - k$ are redundant. Code rates that are specified in the IEEE 802.16 documents include 1/2, 2/3, 3/4 and 5/6.

There are two main categories of FEC: block coding and convolutional coding. The difference is that block codes work on fixed-size blocks while convolutional codes work on symbol streams of arbitrary length.

The mandatory channel coding scheme in IEEE 802.16e-2005 is based on binary non-recursive convolutional coding (CC) with tail biting. It is defined in section 8.4.9.2.1 in the standard and is also mandatory in the mobile WiMAX system profile for both the BS and the MS. In addition, several optional channel coding schemes are defined in the standard, such as zero tail convolutional codes and different turbo codes [7, 8]. Turbo codes are one of the few FEC coding schemes that approach the Shannon limit [18], the theoretical limit of the maximum information transfer rate over a noisy channel.

Convolutional turbo codes (CTC) are worth mentioning because of their superior performance and high popularity in other broadband wireless systems such as HSDPA, WCDMA, and 1xEV-DO [3]. While these systems use CTC with binary turbo encoders, WiMAX specifies CTC with duo-binary turbo encoders, which have several advantages with regards to minimum distance between codewords, less sensitivity to puncturing patterns, and better robustness of the decoder [19]. CTC is listed as mandatory for both the BS and the MS in the mobile WiMAX system profile.

Block turbo codes (BTC) and low density parity check (LDPC) codes are also defined. However, they seem unlikely to be implemented as most equipment manufacturers have decided to go for CTC because of its superior performance over the other schemes [3]. Both are listed as optional in the mobile WiMAX system profile.

2.5.3 Interleaving

The next step after channel coding is bit interleaving, which is used to protect the transmission against long sequences of consecutive errors. The encoded bits are interleaved using a two-step process. The first step ensures that adjacent coded bits are mapped onto non-adjacent sub-carriers. This provides frequency diversity and improves the performance of the decoder. The second step ensures that adjacent bits are alternately mapped to less and more significant bits of the modulation constellation. This is very critical, since for 16-QAM and 64-QAM constellations, the probability of error is not the same for all the bits [3, 12].

Interleaving is defined in section 8.4.9.3 of the IEEE 802.16 standard and is mandatory to implement for the mobile WiMAX system profile if CTC is not used. If convolutional coding is used, there is an optional interleaver that may be implemented.

2.5.4 Repetition

Repetition is a new feature that was added by the IEEE 802.16e-2005 amendment. After channel coding and bit interleaving, the data is segmented into slots. Each slot may then be repeated R times where R is a number from the set $\{2, 4, 6\}$. This scheme applies only to QPSK modulation, and its purpose is to repeat the data enough times to ensure successful reception even at cell edges where the radio link is poor. For instance, the FCH uses repetition coding because it contains information vital to all subscribers. Repetition is defined in section 8.4.9.5 of the IEEE 802.16 standard and is mandatory to implement in the mobile WiMAX system profile.

2.5.5 Modulation

Modulation in a telecommunication system is the process of varying the physical characteristics of a periodic waveform in order to transfer a digital bit stream over an analog channel. The waveform is usually a high-frequency sinusoidal carrier. By adjusting the frequency, phase, amplitude, or a combination of these, different modulation schemes are obtained.

Four digital modulations are supported by the IEEE 802.16 suite of standards, as defined in section 8.4.9.4: BPSK, QPSK, 16-QAM and 64-QAM. In the mobile WiMAX system profile, QPSK, 16-QAM and 64-QAM are listed as mandatory constellations in the downlink, while only QPSK and 16-QAM are listed as mandatory in the uplink. The following sub-sections briefly introduce each of them.

BPSK

Binary phase-shift keying (BPSK) belongs to the family of phase-shift keying (PSK) modulation schemes. They convey data by changing the phase of the carrier wave. BPSK is the simplest form of PSK, and it uses two phases which are separated by 180 degrees. As can be seen from the constellation diagram[1] in figure 2.4, BPSK is only able to encode one bit per symbol. Hence, it is not suitable for applications that require high data rates. However, it has high

[1] A constellation diagram is a representation of a digital modulation scheme in the complex plane. The real and imaginary axes are often called the in-phase axis and quadrature axis respectively.

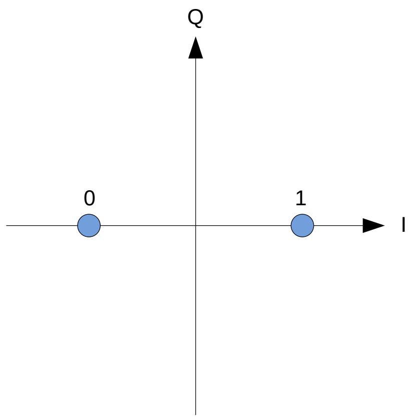

Figure 2.4: BPSK constellation diagram

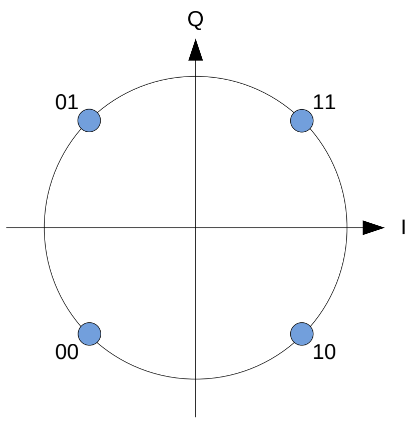

Figure 2.5: QPSK constellation diagram

immunity against noise and interference because it takes serious distortion to wrongly interpret the phase [12].

QPSK

Quadrature phase-shift keying (QPSK) uses four points on the constellation diagram, as shown in figure 2.5. Because it can encode two bits per symbol it doubles the data rate compared to BPSK while maintaining the same bandwidth of the signal. Twice the power is required to achieve the same bit error rate (BER) as BPSK, as two bits are transmitted simultaneously.

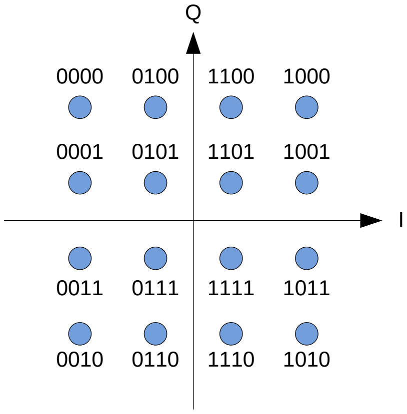

Figure 2.6: 16-QAM constellation diagram

QAM

Quadrature amplitude modulation (QAM) conveys data by changing the amplitude of two carrier waves that are out of phase with each other by 90 degrees. The number of points on the constellation diagram is usually a power of 2, where the most common forms are 16-QAM, 64-QAM, 128-QAM and 256-QAM. Moving to a higher-order constellation allows for the transmission of more bits per symbol, but at the same time the bit error rate increases as the system becomes more prone to noise and interference. In general, a higher-order QAM can deliver more data less reliably.

The constellation diagram of 16-QAM is shown in figure 2.6. A Gray coded bit-assignment is given [20], where adjacent symbols differ in only one digit. This coding reduces the bit error rate (BER) as the most probable error from one symbol to the next produces only a single bit-error.

Adaptive Modulation and Coding

IEEE 802.16 specifies a number of digital modulation schemes. Having more than one scheme is a great advantage when it comes to link adaptation in fluctuating channels. When the radio link is good, a high-level modulation technique and code rate are used that can transfer many bits per symbol time. For example, correctly interpreting 256-QAM requires the carrier to be at least 30 dB, or 1000 times stronger than the noise in the channel [21]. When the radio link is bad, a lower-level modulation technique is used which reduces the data-rate but increases the robustness of the transmission.

	DL	UL
Modulation	QPSK, 16QAM, 64QAM	QPSK, 16QAM, 64QAM
Code rate CC	1/2, 2/3, 3/4, 5/6	1/2, 2/3, 5/6
Code rate CTC	1/2, 2/3, 3/4, 5/6	1/2, 2/3, 5/6
Repetition	x2, x4, x6	x2, x4, x6

Table 2.3: Supported modulations and code rates in mobile WiMAX [14]

Most modern wireless communications systems use adaptive modulation and coding (AMC) techniques that change the modulation, coding rate, and other signal parameters dynamically as the radio link conditions change [12]. Table 2.3 shows the modulation and code rates that are supported by the mobile WiMAX system profile. Note that not all of these are listed as mandatory.

The next step in the chain is space/time encoding if transmit diversity or MIMO functionality is implemented. These techniques belong to the advanced set of features and are dealt with later, in section 2.7.2. Following the space/time encoder, the next set of functions are related to the construction of the OFDM symbols.

First, a general introduction to OFDM is given as this multiplexing scheme is one of the basic building blocks of WiMAX. Following this, the more advanced OFDMA is presented along with its scalable version, which forms the foundation of the mobile WiMAX system profile.

2.5.6 OFDM

Orthogonal frequency-division multiplexing (OFDM) is a digital multi-carrier modulation technique that was proposed by Bell Labs in 1966 [22]. It is used in many modern communication systems, such as ADSL, DVB-T, Flash-OFDM, LTE, 802.11 and WiMAX.

OFDM has several advantages over single-carrier modulation schemes. The primary advantage is its ability to cope with severe NLOS environments by effectively overcoming inter-symbol interference (ISI) and frequency-selective fading caused by multipath propagation[2]. It has very high spectral efficiency, near the Nyquist rate, and facilitates the use of single frequency networks where signals from multiple distant transmitters may be combined constructively.

[2]Multipath is the reception of multiple versions of the same signal, each which has travelled a different path between the transmitter and receiver.

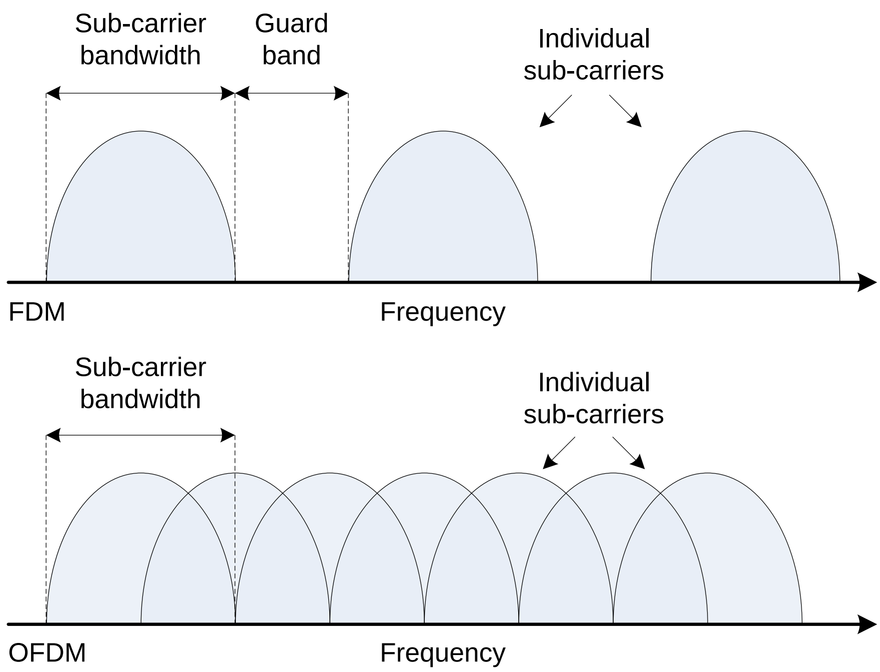

Figure 2.7: FDM versus OFDM

By definition, a high data-rate system like WiMAX operating in a NLOS environment will generally have a channel delay spread[3] τ much larger than the symbol time T_s, since the number of symbols sent per time unit is high [3]. If the delay spread becomes larger than the symbol time, ISI is introduced and the communication system fail as the bit error rate becomes intolerable. In order to overcome this problem, the symbol time must be made significantly larger than the channel delay spread. Multi-carrier modulation techniques such as OFDM solves this by dividing the high bit-rate data stream into L parallel lower bit-rate streams called sub-carriers, *each* of which has $T_s/L \gg \tau$. Hence, in theory the sub-carriers become effectively ISI free and the total data rate is maintained. In order to create such a large number of sub-carriers without the need for L radio frequency (RF) radios in both the transmitter and receiver, OFDM uses a very efficient implementation of the discrete Fourier transform (DFT) known as the fast Fourier transform (FFT). Both the FFT and its inverse, the IFFT, are used to combine and create a multitude of orthogonal sub-carriers using a single radio.

[3]The channel delay spread specifies the duration of the channel impulse response $h(\tau, t)$, which is the time between the first arriving signal and the last non-negligible signal at the receiver. This phenomenon is caused by multipath propagation.

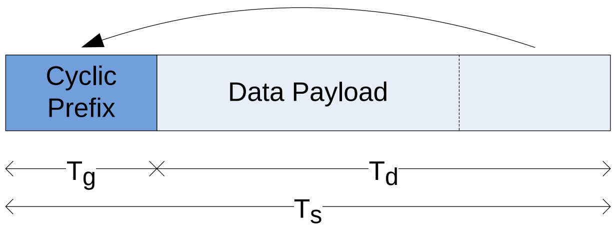

Figure 2.8: Cyclic prefix [14]

Each sub-carrier is then modulated with a conventional modulation technique such as BPSK, QPSK, or QAM. These sub-carriers typically overlap in frequency, but are designed in such a way that they are orthogonal to each other over the duration of the symbol. This means that the peak of one sub-carrier coincides with the nulls of the adjacent sub-carriers. Because of the orthogonality, inter-carrier interference[4] (ICI) is eliminated and guard bands between the sub-carriers are not required. The latter greatly improves the spectral efficiency of OFDM compared to i.e. frequency division multiplexing (FDM), where a guard band is needed to separate each of the non-overlapping sub-carriers. This principle is illustrated in figure 2.7. Also, the design of both the transmitter and receiver becomes more simple with OFDM compared to FDM for high data rates, as separate filters for each sub-carrier is not required [3, 12].

Even though there are no need for guard bands between individual sub-carriers because of their orthogonality, OFDM still requires guard bands between adjacent symbols. This is necessary in order to keep each of the symbols independent of the others after going through the wireless channel. The length of the guard interval is made large enough so that the channel-induced delays are an insignificant fraction of the total symbol duration. When the guard interval is larger than the expected multipath delay spread, ISI between subsequent OFDM symbols is in theory completely eliminated. However, a large guard interval implies power wastage and decreases the spectral efficiency, hence it is necessary to find an optimal balance [3, 12].

The guard band mitigates ISI between subsequent OFDM symbols, making each symbol independent of those coming before and after it. Interference within an OFDM symbol is mitigated by the use of a cyclic prefix (CP), which is a requirement for the IFFT/FFT operation. The CP is typically a repetition of the last samples of data portion of the block, making the linear

[4] Cross-talk between sub-carriers

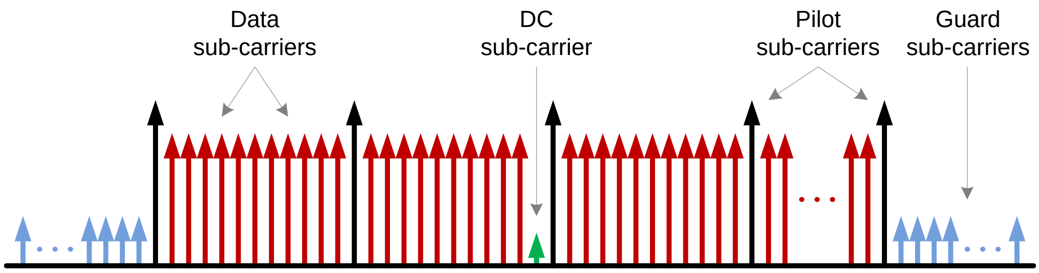

Figure 2.9: Frequency-domain representation of an OFDM symbol [7]

convolution of the channel appear circular. As shown in figure 2.8, the CP is appended to the beginning of the payload data, and occupies a duration called T_g. The duration of the useful symbol time is denoted T_d. The ratio T_g/T_d is often denoted G in IEEE 802.16 documents and can be tuned to suit varying channel conditions [7]. If the radio channel is bad due to multipath, a high value of G is needed. If the radio channel is good, a lower value of G may be used.

For the OFDM and OFDMA PHY layers, IEEE 802.16 defines the following values for G: 1/4, 1/8, 1/16, and 1/32 [7, 8]. However, the mobile WiMAX system profile only specifies 1/8 as mandatory, the others are optional.

All sub-carriers of an OFDM symbol in WiMAX do not carry useful data. As shown in figure 2.9, there are four sub-carrier types that are defined in the IEEE 802.16 documents [3, 7, 8]:

- Data sub-carriers used for data transmission.

- Pilot sub-carriers used for channel estimation and synchronisation.

- Null sub-carriers used for frequency guard bands, which carries no transmission at all. No power is assigned to the guard sub-carriers in order to fit the OFDM symbol within its allocated bandwidth and thus reduce interference between adjacent channels. This is known as spectral roll-off.

- The DC sub-carrier is also a null sub-carrier whose frequency is equal to the RF centre frequency of the transmitting station. The DC sub-carrier is not modulated to prevent saturation effects or excess power draw at the amplifier.

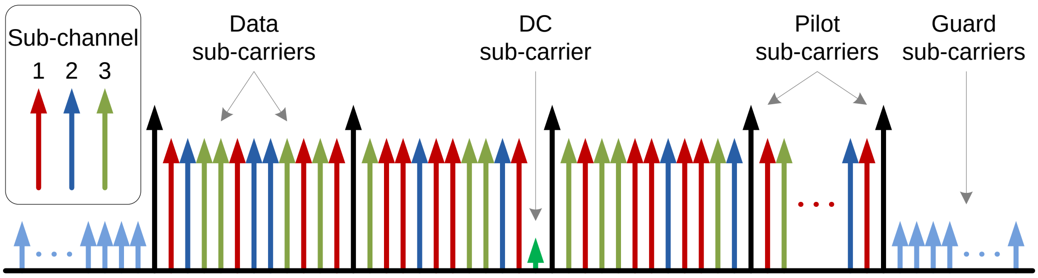

Figure 2.10: Frequency-domain representation of an OFDMA symbol [8]

2.5.7 OFDMA

OFDM is not a multiple-access scheme but rather a modulation technique that creates a number of independent sub-carriers that can be used by different users. Communication systems such as DSL, 802.11a/g and fixed WiMAX use single-user OFDM. This means that all available sub-carriers are used by a single user at a time. In order for multiple users to use such systems concurrently, the available bandwidth is commonly divided through the use of frequency, time, or code division multiple access. In frequency division multiple access (FDMA), each user receives a unique sub-carrier. In time division multiple access (TDMA), each user is given a unique time slot, and in code division multiple access (CDMA), orthogonal binary codes separate one user from the others.

Orthogonal frequency division multiple access (OFDMA) is essentially a hybrid of FDMA and TDMA, where users are dynamically assigned unique sub-carriers in different time slots. In other words, they are separated in both the time and frequency domains. This principle is illustrated in figure 2.10, where the available sub-carriers are grouped into three sub-channels[5], each used by a unique user. As can be seen from the figure, OFDMA allows transmissions from several users at the same time, where each user is only given control over a fraction of the available bandwidth. This flexible multiple-access technique may accommodate users with varying requirements with regards to data rate and QoS.

Another significant advantage of OFDMA over OFDM is its potential to reduce the transmit power and to relax the peak-to-average power ratio (PAPR) problem, which is described as follows. A multi-carrier signal is

[5]The grouping of sub-carriers into sub-channels is known as sub-channelisation, and is dealt with in section 2.5.8

the sum of many narrow-band signals, and at some time instances the sum is large and at others it is small. This means that the peak value of the OFDM signal is substantially larger than the average single-carrier signal. This high PAPR is an important challenge for OFDM systems, especially in the uplink, because it reduces the efficiency, causes difficulties, and increases the cost of the RF power amplifier [3]. In OFDMA, the entire bandwidth is split among several users in the cell, where each user transmits only on a subset of the available sub-carriers. Because PAPR increases with the number of sub-carriers, this causes each user to transmit with a lower PAPR and much lower power than if it had to transmit over the entire bandwidth. The result is that lower data rates and bursty traffic are handled more effectively in OFDMA than in single-user OFDM, since they are sent over a longer time period using the same total power.

2.5.8 Sub-channelisation

A sub-channel in OFDMA is a logical collection of data and pilot sub-carriers. The sub-carriers that make up a sub-channel may be adjacent to each other or distributed throughout the frequency band, depending on the permutation mode [3, 12, 14].

Sub-channelisation was included in the IEEE 802.16-2004 standard for the uplink only. In the IEEE 802.16e-2005 amendment, this was extended to the downlink as well. In the following, the different sub-carrier permutation schemes that are allowed are discussed.

Distributed Permutation Mode

In the distributed permutation mode, the sub-carriers are drawn pseudo-randomly to form a sub-channel. This method provides better frequency diversity and provides inter-cell interference averaging. Permutations include downlink full usage of sub-carriers (DL FUSC), downlink partial usage of sub-carriers (DL PUSC), and uplink partial usage of sub-carriers (UL PUSC).

In DL FUSC, each sub-channel is made up of 48 data sub-carriers, which are distributed evenly throughout the entire frequency band. The set of pilot sub-carriers are divided into two constant sets and two variable sets. The pilots in the variable sets change indices from one OFDM symbol to the next, while the pilots in the constant sets remain unchanged. The variable sets of pilots allow the receiver to more accurately estimate the channel response. This procedure is detailed in the IEEE 802.16 standard section 8.4.6.1.2.2.1

In DL PUSC, all sub-carriers are first arranged into clusters, where each cluster consists of 14 adjacent sub-carriers over two OFDM symbols. These 28 sub-carriers are then divided into 24 data and 4 pilot sub-carriers. The clusters are then renumbered using a pseudo-random numbering scheme, which redistributes the clusters logically throughout the spectrum. After the rearranging scheme, the clusters are divided into six groups. A sub-channel is then created using two clusters from the same group. This procedure is detailed in the IEEE 802.16 standard section 8.4.6.1.2.1.1.

The UL PUSC is analogous to the DL PUSC, where the sub-carriers are first divided into tiles. Each tile consists of four sub-carriers over three OFDM symbols. Six tiles are chosen from the entire spectrum by means of a permutation scheme and grouped together to form a slot. This procedure is detailed in the IEEE 802.16 standard section 8.4.6.2.1.

All of the three aforementioned permutation modes are mandatory for both the BS and MS in the mobile WiMAX system profile.

Adjacent Permutation Mode

In the adjacent permutation mode, a block of contiguous sub-carriers are drawn to form a sub-channel. This method is more desirable for beamforming and multi-user diversity. Permutations include DL adaptive modulation and coding (AMC) and UL AMC.

In this permutation method, nine adjacent sub-carriers are used to form a bin, and four adjacent bins constitute a band. An AMC sub-channel consists of six contiguous bins from within the same band. This can be accomplished by using one bin over six consecutive symbols, two consecutive bins over three consecutive symbols, or three consecutive bins over two consecutive symbols [3]. In the mobile WiMAX system profile, only the 2x3 AMC permutation method is mandatory.

2.5.9 Scalable OFDMA

Operators around the globe may deploy WiMAX networks using different bandwidths of operation, depending on the spectrum allocation scheme in that country. In fixed WiMAX, where the number of sub-carriers is set to

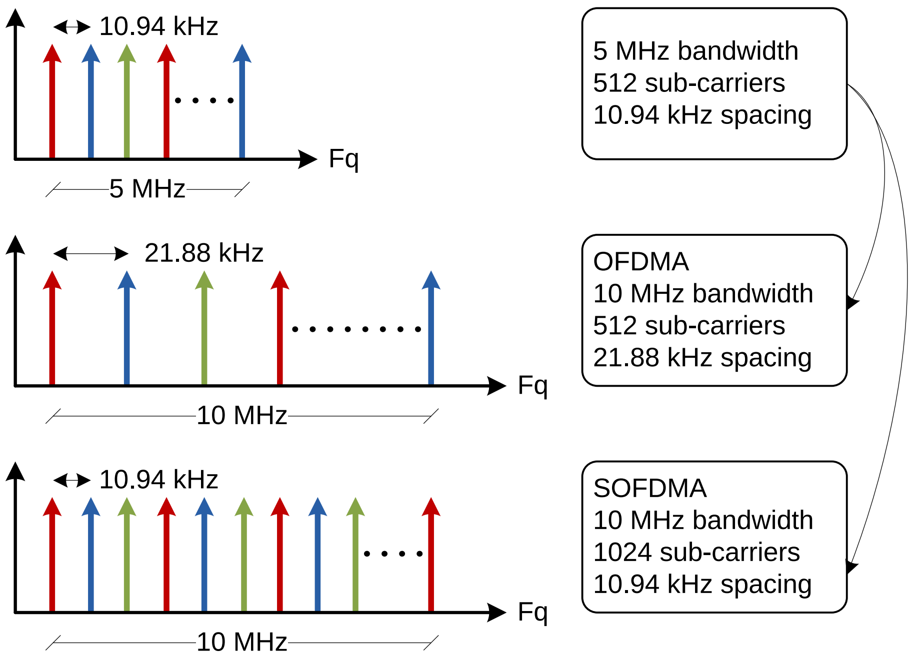

Figure 2.11: Sub-carrier spacing in SOFDMA

256, this requires the sub-carrier spacing and the OFDM symbol time to vary from deployment to deployment. This would not cause problems for fixed operation in a single network, but it would cause trouble for mobile clients moving from one bandwidth of operation to another. One of the problems of operating with different sub-carrier spacing is related to Doppler spread and frequency offsets, where algorithms that are used to combat such issues must be tuned specifically for different bandwidths.

To overcome this problem, the IEEE 802.16e-2005 amendment proposes a technique known as scalable OFDMA (SOFDMA). While the word scalable does not appear in the standard, the technique works by scaling the FFT size and thus the number of sub-carriers. The result is that a fixed sub-carrier spacing is achieved, independent of the bandwidth in use [12, 14]. This concept is illustrated in figure 2.11. As can be seen, the topmost frequency domain illustrates a 5 MHz channel with 512 sub-carriers. If the bandwidth is increased to 10 MHz and the number of sub-carriers remain fixed, the spacing between them doubles, as illustrated in the middle frequency domain. Last, the case with scalable OFDMA is illustrated, where the number of sub-carriers double while keeping their initial spacing distance constant.

The supported FFT sizes are 2048, 1024, 512, and 128, while only 1024 and

512 are mandatory for the mobile WiMAX system profile.

2.5.10 Duplexing

When a communication system is two-way, it is said to be duplex. This means that the transmitter and receiver are able to communicate both ways. There are two forms of duplexing available. frequency division duplexing (FDD) and time division duplexing (TDD).

Frequency Division Duplexing

FDD is achieved by using two separate frequencies for the uplink and downlink. Transmission from one entity to the other is transmitted on one frequency while the return path uses the other frequency. Because such systems usually share a common antenna, the two frequencies must have a large separation between them in order to assure that the locally transmitted energy can be filtered out of the local receiver [21]. FDD is more efficient in systems that experience symmetric traffic because the two channels usually operate over the same bandwidth.

Time Division Duplexing

TDD uses a single frequency for both uplink and downlink. Duplexing is achieved by dividing the air-time into transmit and receive timeslots with a small guard in between them. TDD is useful in systems that experience asymmetric traffic because the ratio between uplink and downlink air-time may be dynamically allocated as traffic conditions change.

The IEEE 802.16e-2005 document specifies both TDD and FDD modes. However, the initial release of mobile WiMAX certification profiles only include support for the TDD mode of operation. The WiMAX Forum states that FDD profiles will be considered to address specific market opportunities where regulatory requirements prohibit the use of TDD [14].

There are several advantages to TDD over FDD when it comes to duplexing mode, some which are listed below [14]:

- TDD enables dynamic allocation of DL and UL bandwidth to efficiently support asymmetric traffic. FDD has a fixed and usually equal bandwidth for both DL and UL, which wastes bandwidth if the traffic is not symmetric.

- TDD ensures channel reciprocity for better support of link adaptation and advanced antenna techniques such as beamforming and MIMO.
- FDD requires a pair of channels while TDD requires only a single channel for both downlink and uplink. This provides greater flexibility for adaptation to varied global spectrum allocations.
- The transceiver design for TDD is less complex and less expensive than for FDD. However, base stations tends to be more complex as they must be synchronised.

Half-duplex FDD

An additional duplexing mode known as H-FDD is specified in the IEEE 802.16e-2005 amendment. It is basically a FDD duplexing scheme where the user cannot transmit and receive at the same time. From a cost and implementation perspective, H-FDD is cheaper and less complex than its FDD counterpart, but the peak data rates are less [3].

The next section presents the frame structure used for the OFDMA physical layer with TDD.

2.5.11 OFDMA Frame Structure

An OFDMA physical layer frame is illustrated in figure 2.12. This frame includes features that are not mandatory in the IEEE 802.16e-2005 amendment, but mandatory in the mobile WiMAX system profile. It differs considerably from the OFDM physical layer frame because it maps data in two dimensions, both time and frequency.

As can be seen from the figure, each frame is divided into two sub-frames, one for DL and one for UL. The two sub-frames are separated by two transmit/receive and receive/transmit transition gaps (TTG and RTG respectively).

The first OFDM symbol in the DL sub-frame is the preamble, which is used for physical layer procedures such as time and frequency synchronisation, initial channel estimation, and noise and interference estimation. Following the preamble, the FCH carries control information about the frame and its

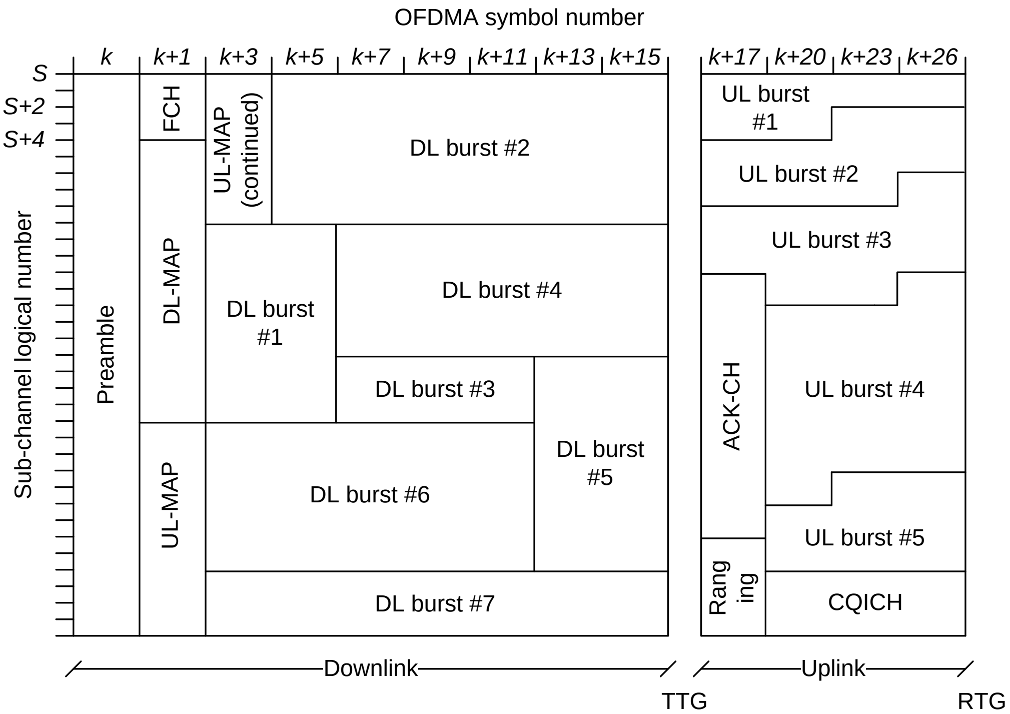

Figure 2.12: OFDMA physical layer frame structure [14]

usable sub-channels [3]. The FCH is always coded with BPSK 1/2 to ensure maximum robustness and reception even at the edge of the cell.

The next two segments are the DL-MAP and UL-MAP. They specify which users are allocated which data regions in the current frame. The rest of the DL frame is made up of several bursts, where each burst may use modulation and coding independent of the other bursts. For instance, a user with good radio conditions may receive data with a high order of modulation, and thus high bit rate, while another user at the cell edge may receive data with more robust coding.

In the uplink sub-frame, the ranging channel is allocated for mobile stations to perform ranging and bandwidth requests [14]. The channel quality information channel (CQICH) is used for the mobile stations to feedback channel state information used for handover and MIMO operations, and last the ACK channel is used by the mobile stations to feed HARQ acknowledgements back to the transmitter, as discussed later in this thesis.

Additionally, the BS periodically transmits channel descriptors which contain additional control information such as channel structure and the various burst profiles that are supported. They are not transmitted in every frame.

2.6 MAC layer

The media access control (MAC) layer of WiMAX is divided into three separate components, as described in figure 2.13. The service-specific convergence sub-layer (CS) at the top interfaces with and receives MAC service data units (MSDUs) from the layer above. It is responsible for performing address mapping and packet header suppression (PHS), where the repetitive part of higher layer headers are effectively removed. This sub-layer is designed to interact with a variety of higher layer protocols, such as ATM, IP, Ethernet, and any unknown future protocol [3]. Support for both IPv4 and IPv6 are required in the mobile WiMAX system profile.

The common part sub-layer works independently of the higher layer protocols and performs fragmentation and concatenation of MSDUs into MAC protocol data units (MPDUs). There exists various types of MPDUs, depending on the data they carry. Some contain MSDUs while others contain bandwidth requests or management messages. Each MPDU frame is prefixed with a generic MAC header (GMH) which contains information about the length of the frame, if CRC or sub-headers are present, or if the frame

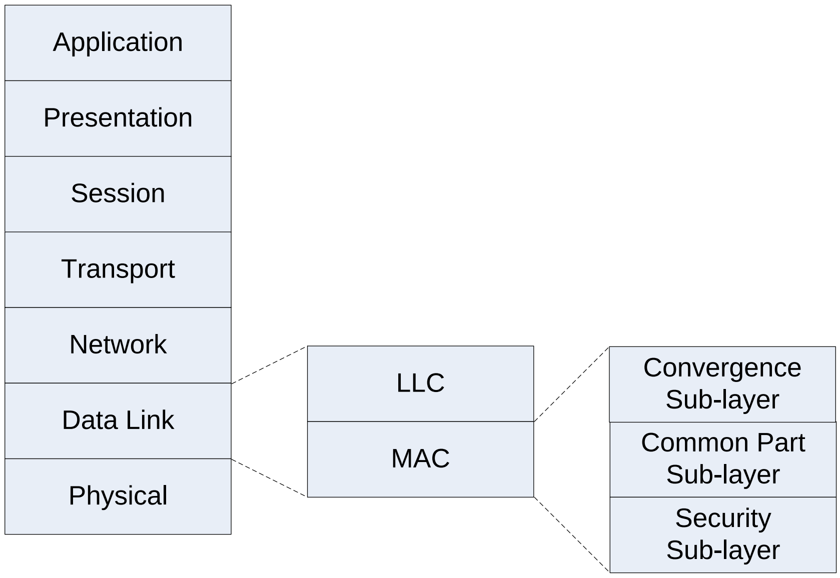

Figure 2.13: The WiMAX MAC layer

is encrypted. The maximum length of a MPDU is 2,047 bytes, represented by 11 bits in the GMH [3]. The common-part sub-layer is also responsible for establishing and maintaining connections with other peers. Further tasks include bandwidth request and allocation, scheduling of packets, modulation and code rate selection, QoS control, and ARQ.

Last, the security sub-layer is responsible for authentication, authorisation, and encryption of user data. It also incorporates a flexible key-management protocol for the exchange of encryption keys between the BS and MS.

In the following, selected aspects of the MAC layer that are of particular importance to mobile WiMAX will be dealt with in more detail. These are QoS, power saving features, mobility management, and security.

2.6.1 QoS

Support for quality of service is a fundamental part of the mobile WiMAX MAC layer, and it uses many techniques and ideas from the well-known DOCSIS cable modem standard [23]. In order to provide strong QoS control, all downlink and uplink connections are managed by the serving BS. Before any data transmission can occur, the BS and MS must establish a unidirectional logical link between the two MAC layer peers [3]. This link is called a connection and it is identified by a connection identifier (CID), which is different for UL and DL. The purpose of the connection is to transport the traffic of a service flow, which is explained next. A connection only exists for one type of service, hence both voice and web-browsing data cannot use the same connection [12].

QoS is provided by the use of service flows, which is a unidirectional flow of packets with a particular set of QoS parameters. Such parameters may be throughput, latency, and jitter requirements. The flow is identified by a service flow identifier (SFID), and various flows with identical QoS parameters are usually grouped into service flow classes.

In order to enforce QoS, outbound MSDUs are first classified according to a set of criteria such as destination IP address, protocol type, DiffServ code points, or MPLS flow labels. They are then mapped onto a particular connection which also creates an association with a service flow that defines its QoS parameters. The classification and mapping mechanisms exist for both uplink and downlink traffic, and they are present in both the BS and the MS. After the classification, the scheduler determines the transmission period for each of the connections, both uplink and downlink. The scheduler is located in the BS, and it also decides what modulation and coding rate to use depending on the radio channel.

The WiMAX MAC layer defines five scheduling services, or QoS classes, to handle traffic with different QoS requirements [3, 12]. Four scheduling services were defined in the IEEE 802.16-2004 standard, while a fifth was added in the IEEE 802.16e-2005 amendment. All of them are mandatory to implement in the mobile WiMAX system profile:

- Unsolicited grant service (UGS) is designed to support real-time service flows that generate fixed-size packets at periodic intervals. Examples of applications are T1/E1 emulation or voice over IP (VoIP) without silence suppression.

- Real-time polling service (rtPS) is designed to support real-time service flows that generate variable-size packets at periodic intervals. Example of application is MPEG video streaming.

- Extended real-time polling service (ertPS) was introduced with the IEEE 802.16e-2005 amendment and is designed to accommodate data services whose bandwidth requirements change with time. Example of application is VoIP with silence suppression.

- Non-real-time polling services (nrtPS) is designed to support delay-tolerant data streams that generate variable-size packets and requires a minimum data rate. Example of application is FTP.

- Best-effort service (BE) provides very little QoS and is applicable for applications which have no strict QoS requirements. Data is sent whenever resources are available, and applications that make use of this class may be web browsing or e-mail services.

2.6.2 Power Saving Features

For the MS to operate for longer periods of time without having its battery recharged, it is critical that power saving features are implemented. Power saving is achieved by turning off some parts in the MS in a controlled manner when it is not actively transmitting or receiving data. This does not only conserve resources at the MS side. The BS also benefits from such power saving features, as there is less waste of air resources to idle stations. The mobile WiMAX system profile supports two modes of power saving operation, sleep mode and idle mode.

Sleep mode is a state in which the MS turns itself off and becomes unavailable for a pre-determined period of time. The period of absence is negotiated with the BS to which the MS is connected. Each sleep window is followed by a listen window, and the length of each sleep and listen window is dependent on the power saving class. There are three power-saving classes that are defined in the standard. In class 1, the sleep window is exponentially increased from a minimum value to a maximum value. This class is best suited for best-effort or non-real-time traffic. In class 2, the sleep window is of fixed length and is followed by a listen window of fixed length. This class is recommended for UGS services. Last, class 3 allows for a one-time sleep window and is typically used for multicast or management traffic [3]. Only class 1 has been listed as mandatory in the mobile WiMAX system profile.

Idle mode is a mechanism that allows the MS to receive broadcast DL transmissions from the BS without registering itself with the network. This allows for even greater power saving than in sleep mode. Before entering idle mode, the MS is assigned to a paging group by the BS. Periodically, the MS wakes up and checks its current location. If it has moved outside its current paging group, it performs a paging group update. This is necessary for the network to know the approximate area where the MS is located. Idle mode is mandatory in the mobile WiMAX system profile.

2.6.3 Mobility Management

The IEEE 802.16e-2005 amendment defines signalling mechanisms for tracking MSs as they move from the coverage of one BS to another when active, or from one paging group to another when idle. There are three methods for handoff that are supported by the mobile WiMAX system profile. One is mandatory while the two others are optional.

The mandatory handoff method is known as hard handover (HHO). In this mode, the MS regularly measures the signal quality of nearby base stations and feeds the information back to its serving BS. Based on these measurements, either the BS or the MS makes a handoff decision, and informs the other peer. Once the decision is made, the MS starts to synchronise with the new BS, performs ranging, and then disconnects from the serving BS before connecting to the new one. This procedure may cause some packets to get lost. In order to optimise the handoff, the WiMAX Forum has developed several techniques with the goal of keeping the delay less than 50 ms [14].

The two optional handoff methods are known as fast base station switching (FBSS) and macro diversity handover (MDHO). In both these methods, the MS maintains a connection with more than one BS simultaneously. With FBSS, the MS maintains a list of possible base stations that it may switch to. This is called the active set. The MS continuously monitors all the BSs in the active set, performs ranging, and maintains a valid connection with each of them. However, the MS only communicates with one of the BSs at any one time, called the anchor BS. When a change of anchor BS is required, the MS signals the desired anchor BS on the CQICH.

In MDHO the MS also maintains a set of BSs, which is called a diversity set. During the handoff procedure, the MS transmits to and receives data from all the base stations in the diversity set simultaneously. In the downlink, multiple copies received by the MS are combined using any of the receive diversity methods, while multiple copies sent to different base stations are aggregated using selection combining. The diversity schemes are introduced in a later section dealing with multiple antenna techniques.

Both FBSS and MDHO require that the base stations in the active or diversity set be synchronised, use the same carrier frequency, and share network entry-related information [3]. This is not fully developed yet and is not part of the mobile WiMAX release 1.0 network specification.

2.6.4 Security

Security is handled at multiple layers of the network, where each layer deals with different aspects of the threat model. In WiMAX, the security sub-layer within the MAC deals with data link layer security such as authentication, authorisation, and encryption.

When first designing the security sub-layer in WiMAX, the working group sought to avoid the mistakes of IEEE 802.11 by incorporating a pre-existing authentication and confidentiality framework into the standard. They implemented parts of DOCSIS SP-BPI+ [23], which was originally designed for cable modems and wired technology. However, because WiMAX is a wireless technology, the two are exposed to different threat models and consequently the security framework somewhat fails to provide adequate protection [24]. For instance, as authentication is not mutual, the system may be prone to a rogue base station attack.

Security requirements for mobile services are also very different from fixed services. Because of the security holes in the IEEE 802.16-2004 standard, the security sub-layer was completely rewritten in the IEEE 802.16e-2005 amendment. New and better encryption algorithms, mutual authentication, support for handoffs, and a new integrity control algorithm were implemented. In the following, the security features that have been listed as mandatory in the mobile WiMAX system profile are briefly mentioned.

The privacy and key management protocol version 2 (PKMv2) forms the basis of mobile WiMAX security. The protocol is used to securely transfer keying material from the base station to the mobile station as well as periodically re-authorising and refreshing the keys [3].

Authentication of users and devices are provided by the extensible authentication protocol (EAP) framework, which supports a wide range of credentials. It is up to the implementer to decide on the authentication method, but the following are commonly used: usernames and passwords, digital certificates, smart cards, subscriber identity module (SIM), or universal SIM (USIM). Authentication is mutual, which means that the BS authenticates the MS and vice versa. This mitigates the risk of rogue base stations. Every user terminal comes with a built-in X.509 certificate which may be used for device authentication.

Encryption of user data is taken care of by AES in CCM mode. CCM is a mode of operation for cryptographic block ciphers which provides both authentication and privacy. It combines the counter mode of encryption with the cipher block chaining message authentication code (CBC-MAC) mode of authentication. The keys used for encryption are generated from the EAP authentication using the PKMv2 protocol.

Control messages are protected from tampering by using the AES-based cipher-based message authentication code (CMAC). This method is mandatory in the mobile WiMAX system profile, while the hash-based message authentication code (HMAC) is listed as optional.

Last, support for fast handover is supported by a three-way handshake that optimises re-authentication with the new BS. This mechanism is specially crafted to prevent any man-in-the-middle attacks.

2.7 Advanced Features

This section introduces some of the advanced features that have been incorporated into the mobile WiMAX system profile to enhance coverage, provide better reliability, and higher throughput.

2.7.1 Hybrid ARQ

In an automatic repeat request (ARQ) protocol, the transmitter adds error detection bits to the data in order for the receiver to detect erroneous frames. When forward error correction bits are added to the ARQ protocol, the schemes are known as HARQ. Note the difference between detecting an error and being able to actually correct it on the fly. Clearly, the latter requires more redundant data and complex algorithms.

The IEEE 802.16e-2005 amendment defines both type I HARQ and type II HARQ. Type I HARQ is the simplest of the two and is referred to as chase combining (CC). If the receiver fails to decode a packet, a retransmission request in the form of a non-acknowledgement (NACK) is fed back to the transmitter. Upon reception of this NACK, the transmitter sends the same coded packet again. The receiver now uses the current and all previous HARQ retransmissions of the data block to decode it. With each new retransmission, the probability of successfully decoding the packet increases. This process continues until the packet is decoded without error or the maximum number of allowable HARQ retransmissions is reached, whereby higher layer protocols initiates retransmission of the data block [3].

Type II HARQ, also known as incremental redundancy (IR), uses a more intelligent method where each retransmission consists of new redundancy bits from the channel encoder. This means that the redundant bits are changed from one transmission to the next. Type II HARQ is more complex than type I HARQ and requires a larger receiver buffer size and also additional signalling.

According to [25, 26], a type II HARQ system may have the potential of achieving better link-level performance compared to a type I HARQ system. The analysis shows that IR over CC coding gains tend to increase with increasing coding rates and higher orders of modulation, while they decrease with lower orders of modulation or with varying signal-to-noise (SNR) conditions. In channels that experience fast time-varying fading or long retransmission delays, the CC scheme may actually outperform IR.

Both types of HARQ are listed as optional in the IEEE 802.16 documents. For the mobile WiMAX system profile, type I HARQ is required for use with CTC while type II HARQ is optional.

2.7.2 Multiple-Antenna Techniques

The use of multiple antennas in wireless communication systems have many advantages. As such, they can be used to increase the coverage area, provide higher reliability of the transmission, or increase the throughput. The mobile WiMAX system profile defines a number of multiple antenna schemes that may be used for such purposes. This sub-section investigates some of these schemes.

The Challenges of Multipath Fading

One of the greatest challenges to traditional wireless systems has been managing multipath fading environments. Fading is caused by the reception of multiple versions of the same signal, each which has travelled along a different path between the transmitter and the receiver. Causes of multipath may be atmospheric ducting, ionospheric reflection and refraction, and reflection from various obstacles such as buildings, trees, hills, or other man-made features. Each signal version will experience differences in attenuation, delay, and phase shift while propagating through the air. At the receiver, because multiple copies of the same signal arrive at different points in time with varying properties, either constructive or destructive interference is experienced.

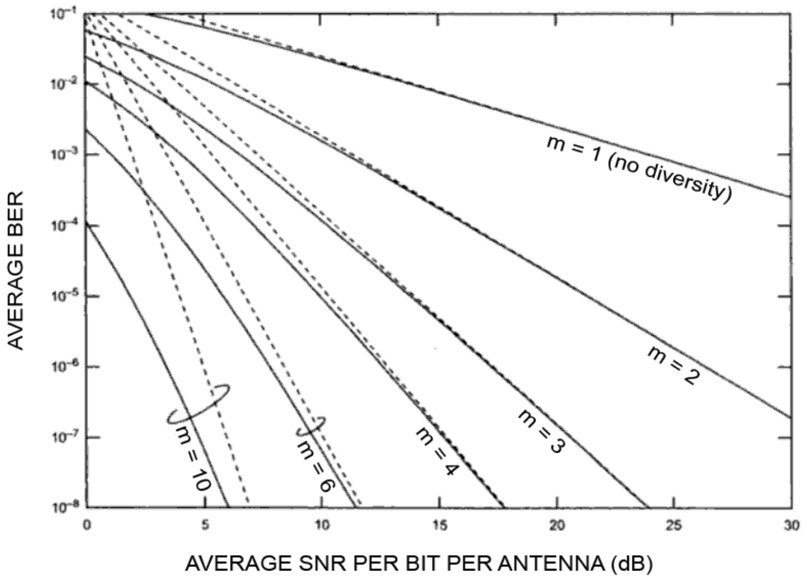

Figure 2.14: Performance of order-m diversity over a Rayleigh-fading channel [29]

This causes the signal to be either amplified or attenuated, depending on the fading that was experienced in the channel [3, 27].

Several techniques are used to correct or mitigate the effects of multipath fading. One common method is to transmit the signal over multiple channels that each experience independent fading. At the receiver, these independent channels are then combined in some fashion, making the probability of experiencing deep fades on all channels simultaneously very low [28]. This is known as micro diversity, and its main application is to improve the relationship between the BER and SNR.

Generally speaking, in a Rayleigh-fading channel with no diversity, the average BER is inversely proportional to the SNR. This relationship is mathematically defined as $BER \sim SNR^{-1}$, which implies that the average BER decreases very slowly as the SNR increases. This fact is illustrated in figure 2.14, with the line labelled no diversity. When order-m diversity is introduced, the relationship improves to $BER \sim SNR^{-m}$ for high values of SNR. In other words, on a log-log scale, the number of diversity branches m defines the slope of the BER versus SNR curve. A higher diversity order leads to a steeper slope [29]. It is also interesting to note that the improvement in BER from order-1 to order-2 diversity is greater than from order-2 to order-3.

There are several ways to achieve diversity, however they are commonly separated into time, frequency, or space:

- Time diversity makes use of coding/interleaving or adaptive modulation techniques. Coding and interleaving introduces redundancy in the transmitted signal causing each symbol to have its information spread over a number of channel coherence times[6]. Adaptive modulation dynamically chooses a modulation technique that achieves the highest possible data rate while still meeting its BER requirement.

- Frequency diversity transmits the signal using several frequency channels, or spread over a wide spectrum, in order to have its information spread over different channel coherence bandwidths[7]. Commonly used techniques to mitigate multipath include OFDM, spread spectrum such as direct sequence and frequency hopping, and the use of Rake receivers.

- Spatial diversity is achieved by using two or more antennas at the receiver side and/or transmitter side. The simplest form is known as selection combining, where the receiver utilises two antennas and dynamically selects the stronger of the two signals received. More sophisticated forms include receive antenna arrays with maximum receive ratio combining, transmit diversity using space time codes, and combinations of these.

For the purpose of this thesis, the spatial diversity schemes are of great importance. In the next sub-sections, this scheme will be dealt with in more detail.

Spatial Diversity

Spatial diversity is achieved in space by transmitting the signal over several different propagation paths. This can be implemented by using multiple transmit antennas (transmit diversity) and/or multiple receive antennas (receive diversity). Additionally, antenna arrays can be used to focus the energy (known as beamforming), or create multiple parallel channels for carrying unique data streams (spatial multiplexing) [3]. When using multiple antennas at both the transmitter and the receiver, these approaches are often

[6] The coherence time is the time over which the channel may be considered virtually constant.

[7] The coherence bandwidth is a statistical measurement of the bandwidth over which the channel transfer function remains virtually constant.

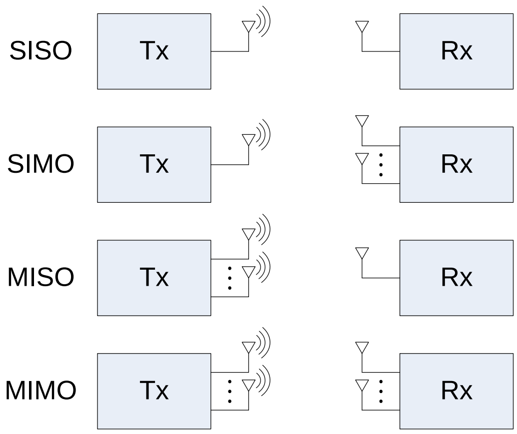

Figure 2.15: Different MIMO schemes [30]

collectively referred to as multiple-input multiple-output (MIMO) and can be used to:

1. Increase the system reliability
2. Increase the achievable data rate and hence system capacity
3. Increase the coverage area
4. Decrease the required transmit power

The term MIMO generally assumes multiple antennas at both the transmitter and receiver. However, spatial diversity is also possible using a number of antenna configurations, including the use of a single antenna at one side of the communication link. Figure 2.15 illustrates the four commonly used scenarios, where SISO, SIMO, and MISO are all degenerate cases of MIMO. Note that the input and output relate to the medium, which explains why multiple input means multiple antennas at the transmitter side sending information into the air.

While the cost associated with additional antenna elements and their accompanying RF chains is not negligible, the gain from such antenna arrays is so enormous that there is little question that emerging 4G mobile broadband standards like WiMAX and LTE will embrace MIMO [3]. Early deployments will likely be conservative in the number of antennas deployed, with primary focus on equipping the base station. In time, when the cost of RF components in subscriber units drop, multiple antennas will be introduced in this segment also.

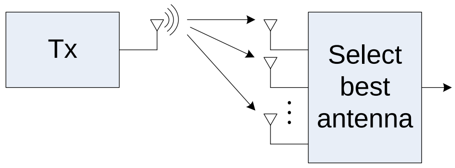

Figure 2.16: Selection combining receive diversity [3]

In the following, several multiple-antenna techniques are listed. The first section introduces receive diversity, which is the most well-established form of spatial diversity. The next section deals with transmit diversity, which requires a different approach. Following this, beamforming and spatial multiplexing are discussed.

Receive Diversity

Receive diversity is the most commonly used form of spatial diversity, and it is nearly ubiquitous in cellular base stations and wireless LAN access points. Its main advantage is that it places no particular requirements on the transmitter. The receiver, however, must be able to process and combine a number of independently received streams in some fashion [3]. Because receive diversity requires multiple antennas and RF chains at the receiver, the cost, complexity, and power consumption increases. As a result, this technique has almost exclusively been applied to base stations to improve their reception capability. The base stations serve a number of subscribers, hence it is more economically feasible to add receive diversity equipment here, rather than at all the subscriber units. However, as prices drop and equipment becomes smaller and easier to implement, receive diversity functionality is likely to be added to the subscriber units as well.

Two widely used algorithms are selection combining and maximum receive ratio combining.

Selection combining (SC) is the simpler of the two receive algorithms. As can be seen from figure 2.16, the receiver estimates the instantaneous strength of each of the incoming streams and selects the strongest one. One drawback of this method is that it completely ignores the useful energy on the other streams, hence SC is clearly suboptimal.

According to [3], if the N_r received signals are uncorrelated and Rayleigh distributed, the average received SNR for all the branches can be shown to be

$$\begin{aligned}\bar{\gamma}_{SC} &= \bar{\gamma}\sum_{i=1}^{N_r}\frac{1}{i} \qquad (2.1)\\ &= \bar{\gamma}(1+\frac{1}{2}+\frac{1}{3}+\cdots+\frac{1}{N_r})\end{aligned}$$

where $\bar{\gamma}$ is the average received SNR at that location. As can be seen from the equation, each added uncorrelated antenna increases the average SNR, but with rapidly diminishing returns in gain.

Maximum receive ratio combining (MRRC) combines the information from all the received branches in order to maximise the ratio of SNR. As shown in figure 2.17, each branch is weighted by a complex factor q_i that is proportional to the signal amplitude. Branches with better signal energy is enhanced whereas branches with lower energy is attenuated. According to [3], the resulting SNR can be found to be

$$\gamma_{MRRC} = \sum_{i=1}^{N_r}\gamma_i \qquad (2.2)$$

As can be seen, the total SNR is achieved by simply adding up the branch SNRs when the appropriate weighting factors are used. In other words, the SNR growth is linear with the number of receive antennas added. From the Shannon-Hartley theorem it can be observed that since $C = B\log(1+SNR)$, the capacity growth is logarithmic since MRRC linearly increases the SNR [3].

For an interference-limited system such as WiMAX, MRRC would be strongly preferred to SC, despite the fact that the latter technique is somewhat simpler to implement. The two methods have the same diversity gain, but MRRC achieves a higher average SNR. This is because it considers the combined signal on all antennas, whereas SC would simply select the best *average* antenna.

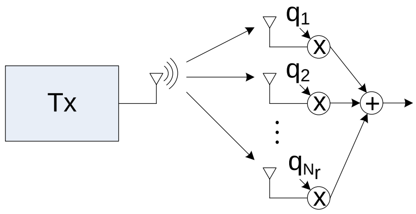

Figure 2.17: Maximum receive ratio combining receive diversity [3]

Transmit Diversity

Transmit diversity schemes improve the signal quality at the receiver side by simple processing across a number of antennas at the transmitter side. In essence, the transmitter sends multiple copies of the data stream across a number of antennas, which the receiver combines in an optimal way. This scheme is particularly attractive for the downlink of infrastructure-based systems because it puts the burden of transmitting from multiple antennas on the base station side, which in many cases is already equipped with a number of antennas used for receive diversity. The incremental cost of using the same antennas for transmit diversity is very low.

Transmit diversity schemes are categorised as either open-loop or closed-loop. Open-loop systems do not require knowledge of the channel at the transmitter side while closed-loop systems do.

One of the most popular open-loop transmit diversity scheme is space/time coding, which was first invented in the early 1990s. There are many types of space/time codes, but the most interesting with regards to WiMAX are cyclic delay diversity (CDD) and space/time block codes (STBCs), which lend themselves to easy implementation. The coding ensures that each transmitted signal is orthogonal to the other, thus reducing self-interference and improving the capability of the receiver to distinguish between multiple signals [27]. Both codes are fundamentally used to enhance system coverage.

In 1998, the Alamouti STBC [31] was suggested. It soon became popular because it was easy to implement and optimal with regards to diversity order. The most simple Alamouti code corresponds to two transmit antennas and a single receive antenna. If two symbols s_1 and s_2 are to be transmitted, the following is sent over two symbol times, where the * denotes complex conjugate:

$$\begin{bmatrix} s_1 & s_2 \\ -s_2^* & s_1^* \end{bmatrix}$$

In the matrix above, each row represents a different time slot, while each column represents a different antenna. This code is referred to as a rate 1 code, as the data rate is neither increased nor decreased. Even though the data is sent twice, the scheme requires no additional bandwidth as the redundancy is applied in space across the multiple antennas, and not in time or frequency [31].

In order to achieve even better results, both MRRC and STBC may be used simultaneously. In this case, the receiver may implement multiple antennas and use receive diversity techniques to combine the signals for more robust reception.

Beamforming

Beamforming is the concept of using multiple antenna elements to adjust the strength of the transmitted or received signals based on the direction, which can be either the physical direction or the direction in a mathematical sense. By applying different weights for each antenna with regards to amplitude and phase, the energy is focused in the direction of the intended receiver or transmitter. Two principal classes of beamforming have been defined: direction of arrival (DOA), which is physically directed, and eigenbeamforming which is mathematically directed [3].

There are many benefits of beamforming. Of the most important are increased coverage, power saving at the MS side, interference mitigation, and capacity increase [12]. By concentrating the receive or transmit energy in the direction of the user, the antenna gain in that direction is greatly increased, and interference from other sources are mitigated. This improves the signal quality and hence the capacity as higher order modulation techniques may be utilised.

The beamforming technique requires a great deal of knowledge about the radio environment to perform well. It is therefore classified as a closed-loop system. Because of the complexity and intelligence required of such a system, it is most often implemented at the BS side. Usually, between two and eight antenna elements and transceivers are used to provide the necessary beamforming capabilities [12].

Further advances in beamforming technology may lead to spatial division multiple access (SDMA). Provided that two users are sufficiently separated in space, it is possible to transmit two separate streams to each user at the same time and on the same physical resources. The two transmitted streams would radiate to each user without interfering with one another. However, SDMA may be difficult to utilise effectively in a mobile environment where the MSs move around and may be separated at one instant but in the same direction at the next instant [12].

Spatial Multiplexing

The basic idea of spatial multiplexing (SM) is to divide a high data-rate stream into N_t independent streams and simultaneously transmit each stream over a different antenna. This requires independent and uncorrelated channels and also very good SNR conditions. If all the streams can be successfully decoded at the receiver, the nominal spectral efficiency is thus increased by a factor of N_t. Theoretically, the capacity increases linearly as a function of $min(N_t, N_r)$, where N_t is the number of transmit antennas and N_r is the number of receive antennas [3, 12].

In order to make use of spatial multiplexing, several antennas and transceivers must be implemented at both the BS and the MS. The full benefit of SM is limited to the lesser of the number of antennas at either side, hence full efficiency requires parity in the number of antennas. On the MS side, this leads to more hardware and processing requirements, along with higher power consumption. Efficient implementation of spatial multiplexing on the MS side has indeed some challenges.

If the MS is only equipped with a single antenna, two or more MSs may join to transmit simultaneously in the same slot. This technique is referred to as uplink collaborative SM and works as if a number data streams were spatially multiplexed from multiple antennas of the same MS [14].

Support in WiMAX

The IEEE 802.16-2004 standard provides some support for multiple-antenna schemes. It specifies the Alamouti STBC and beamforming, but beamforming is not included in the fixed WiMAX system profile.

The IEEE 802.16e-2005 amendment provides extensive support for multiple-antenna schemes. It specifies the use of multiple antennas at both the trans-

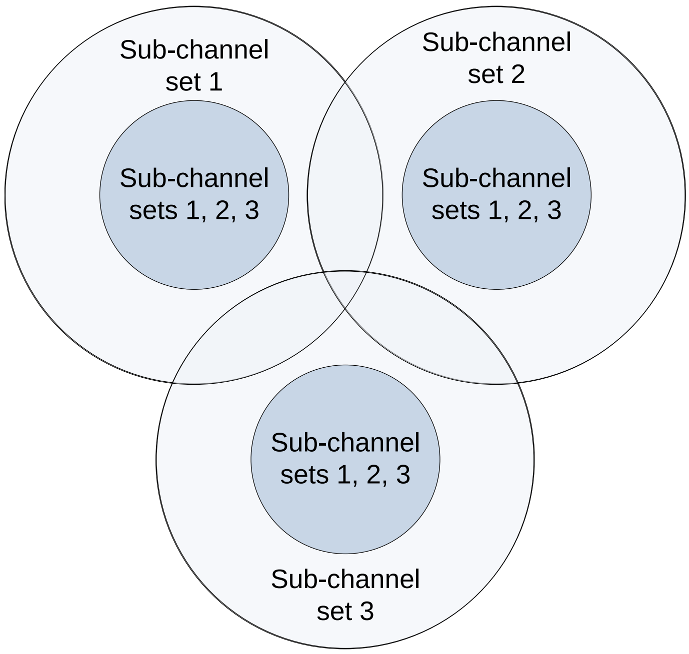

Figure 2.18: Fractional frequency reuse [14]

mitter and receiver for different purposes such as transmit and receive diversity, beamforming and spatial multiplexing. The WiMAX Forum has listed the following schemes as mandatory in the mobile WiMAX system profile: In the DL, STBC and SM is required while only collaborative SM is required in the UL. Note that STBC is referred to as MIMO matrix A and SM as MIMO matrix B in the specifications [14, 17].

2.7.3 Frequency Reuse

WiMAX supports frequency reuse plans where all cells or sectors operate using the same frequency channel to maximise the spectral efficiency. However, users that operate from the edges will experience severe co-channel interference (CCI) from overlapping cells if the full bandwidth is available throughout the entire sector. WiMAX mitigates this problem by allocating a set of sub-channels to each user at the edge such that there is minimal overlap from cell to cell. Because each user only occupies a small fraction of the total bandwidth, the cell edge interference problem can be reduced. This method allows for a more dynamic frequency allocation plan as opposed to traditional fixed frequency planning [3].

The set of sub-channels can be configured in such a way that users under good signal-to-interference ratio (SIR) conditions have access to the full channel bandwidth. These users operate with a frequency reuse of 1. Those users who have poor SIR conditions will be allocated non-overlapping sub-channels which occupy only a fraction of the total bandwidth. These users operate with a frequency reuse of 2 or more, depending on the number of non-overlapping sub-channel groups. This concept is illustrated in figure 2.18.

As discussed earlier, the mobile WiMAX system profile supports a variety of sub-channelisation schemes. All these schemes make it possible to achieve frequency reuse in a flexible manner.

2.7.4 Multicast and Broadcast Service

Multicast and broadcast services (MBS) may be required when multiple MSs connected to the same BS receive the same information. Instead of allocating a dedicated stream to each user, resources may be saved by transmitting identical data to all the users simultaneously. This is of particular interest for broadcast TV types of applications, where several users in the same area receive the same service.

The mobile WiMAX system profile supports downlink MBS by either constructing a separate zone for the service along with unicast data, or by dedicating the whole downlink frame to MBS.

2.8 Network Architecture

The IEEE 802.16 standard specifies only the physical and MAC layer of the radio link between the BS and MS. However, only specifying the air-interface is not sufficient for building an interoperable network that supports multi-vendor access networks, roaming, and inter-company billing. There is a need to specify a multitude of other network functionalities such as session management, end-to-end security, QoS, and mobility management [14].

This task was taken upon by two working groups of the WiMAX Forum. The first is the network working group (NWG), which focuses on creating higher-level network specifications beyond what was defined in the IEEE 802.16 standard. The second is the service provider working group (SPWG),

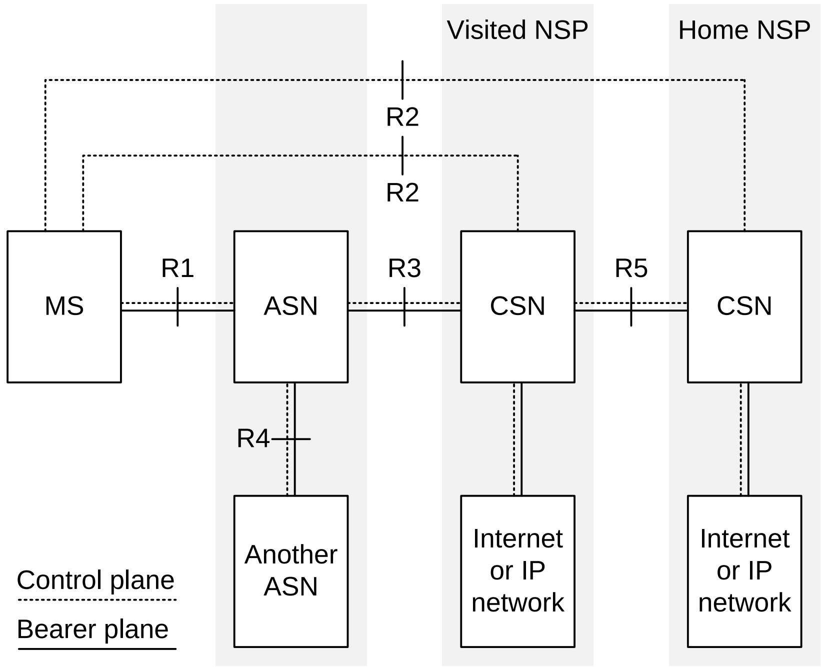

Figure 2.19: Network reference model [14]

which gives service providers a platform for influencing the development of WiMAX.

The end-to-end network system architecture developed by the NWG followed several design tenets [3, 14]. For instance, the architecture should be based on a packet-switched framework, and network-layer procedures and protocols should be based as much as possible on existing IETF protocols. It should be modular and flexible enough to allow for a broad range of implementation and deployment options, from small-scale to large-scale networks. Also, the architecture should support a variety of business models and allow for a logical separation between (1) the network access provider (NAP), which owns and operates the access network, (2) the network service provider (NSP), which owns the subscribers and provides broadband access, and (3) the application service provider (ASP).

Based on these tenets, a network reference model (NRM) has been developed. It specifies the end-to-end network requirements and protocols, and is a logical representation of the network architecture. The model is depicted in figure 2.19 and is logically divided into three parts. (1) Mobile stations used by subscribers to access the network, (2) the access service network (ASN), which is owned by a NAP, and (3) the connectivity service network (CSN), which is owned by an NSP and provides IP connectivity and IP core network functions [3]. The NRM also identifies reference points for interconnection of the three logical parts, labelled R1 to R5.

It is important to realise that each of the three logical parts only represents a grouping of functional entities. Each of these functions may be realised in a single physical device or distributed across multiple devices. This is an implementation choice, and the manufacturer may choose any physical implementation as long as it meets the functional and interoperability requirements [14].

In the following, the functions of the ASN and CSN will be dealt with in more detail [32].

2.8.1 ASN Functions

The ASN shares the R1 reference point with a MS, R3 with a CSN and R4 with another ASN. It may be decomposed into one or more base stations and one or more ASN gateways (ASN-GW).

The BS is defined as a logical entity that implements a full instance of the mobile WiMAX MAC and PHY over the R1 reference point. It represents one sector with one frequency assignment. Additional functions performed by the BS may include scheduling and traffic classification, relaying messages between the MS and ASN-GW, and proxy functionality. The BS may be connected to one or more ASN-GW for load balancing or redundancy purposes.

The ASN-GW provides location management, paging, admission control, and routing of traffic to selected CSNs. It also acts as an authenticator and provides mobility session management. The ASN-GW may be decomposed into two groups of functions, namely the decision point (DP), and the enforcement point (EP). When decomposed in such a way, the DP functions may be shared among several ASN-GWs.

The WiMAX NWG has defined multiple profiles for the ASN. Each of the profiles defines a different decomposition of its functions. For instance, profile B calls for a single entity that combines all the functions of the BS and ASN-GW, while profiles A and C split the functions between the two.

2.8.2 CNS Functions

The CSN is defined as a set of network functions that provides IP connectivity services to the WiMAX subscribers. It may comprise network elements such as routers, AAA servers, firewalls, user databases, and gateways to other networks. The CSN may provide functions such as IP address allocation, Internet access, subscriber billing, QoS management, tunnelling, and roaming.

2.8.3 Mobility Management

Earlier in this thesis, mobility management on the MAC layer was discussed. The focus was on different handoff mechanisms that would allow a MS to switch from one base station to another. Note that those methods were concerned with handoffs over the air-interface only.

In this section, the focus is on mobility for the network as a whole. More specifically, the WiMAX NRM supports two types of mobility. The first is called ASN-anchored mobility, intra-ASN mobility, or micro-mobility. The other is known as CSN-anchored mobility, inter-ASN mobility, or macro-mobility.

ASN-anchored mobility supports handoff scenarios where the MS moves from one base station to another within the same ASN. This type of movement is invisible to the CSN and does not have any impact at the network- or IP-layer levels. Thus, implementing ASN-anchored mobility does not require any changes to the MS [3].

CSN-anchored mobility refers to mobility across different ASNs. Because CSN-anchored mobility involves mobility across different NAPs and IP subnets, IP-layer mobility techniques are required. In WiMAX, mobile IP (MIP) is used to enable such functionality. More specifically, two types of MIP implementations are defined for supporting CSN-anchored mobility. The first one is based on having a MIP client installed on the MS, while the other relies on a proxy MIP in the network.

2.9 Other Wireless Broadband Technologies

WiMAX is not the only solution for delivering wireless broadband. There exists a number of other technologies, both for fixed and mobile operation.

Some of these are proprietary, for instance iBurst from ArrayComm, and Flash-OFDM from Flarion/Qualcomm. Others are standard-based such as 3G cellular systems, IEEE 802.11-based Wi-Fi systems, and WiBro.

In this section, WiMAX is compared and contrasted to some of these technologies.

2.9.1 3G Cellular Systems

Cellular operators around the world are upgrading their networks to support the delivery of broadband applications to their subscribers. There are three 3G standards that are widely used as of today, namely wideband CDMA (WCDMA), CDMA2000, and time division synchronous CDMA (TD-SCDMA). WCDMA is the technology used in the UMTS system while TD-SCDMA is a mobile communication system pursued in the People's Republic of China.

Operators using global system for mobile communications (GSM) are deploying universal mobile telephone system (UMTS) as part of their 3G evolution. Further developments from the 3rd generation partnership project (3GPP) include high-speed downlink packet access (HSDPA) with release 5, high-speed uplink packet access (HSUPA) with release 6 and long term evolution (LTE) with release 8 and onwards.

HSDPA is a downlink only air-interface and is capable of providing a gross throughput of 14.4 Mbps using a 5 MHz channel. However, realising this rate requires use of all 15 codes, which is unlikely to be implemented. In practise, HSDPA supports data rates of 3.6 Mbps and 7.3 Mbps using 5 and 10 codes respectively. In addition to HSDPA, HSUPA provides higher uplink speeds. Throughput of 5.8 Mbps and 11.5 Mbps are standardised as part of release 6 and 7 respectively.

As of writing, there are many cellular networks that implement both HSDPA and HSUPA. Upgrading to HSDPA started in 2006 and continued throughout 2007, while HSUPA was introduced in late 2007. UMTS release 8 has yet to be finalised, but the current state of the standard aims at peak downlink rates of 100 Mbps using 20 MHz of spectrum, yielding a spectral efficiency of 5 bits/sec/Hz.

The 3rd generation partnership project 2 (3GPP2), which is not to be confused with 3GPP, has defined EV-DO as a high-speed data standard. 1xEV-DO was developed by Qualcomm in 1999 and supports peak downlink data rates of 2.4 Mbps in a 1.25 MHz channel. This standard has been upgraded to revision A, which provides up to 3.1 Mbps gross throughput in the downlink and 1.8 Mbps in the uplink channel. A further revision B, also known 3xEV-DO, supports data rates of up to 4.9 Mbps.

In comparison with the aforementioned technologies, WiMAX has several advantages. While the other 3G systems use fixed channel bandwidths, WiMAX allows for flexible deployments using a selectable channel bandwidth. WiMAX also supports dynamic adjustment of the uplink and downlink ratio, which allows for symmetric links such as T1 emulation and a network which adapts dynamically on a frame-to-frame basis to the varying needs of the users.

According to several sources [3, 33, 34], WiMAX can achieve higher spectral efficiency than what is typically achieved in today's 3G systems. For instance, the standard incorporated support for multiple antennas right from the start, while 3G systems needed to add this feature in the form of revisions. Moreover, exploiting multi-user and frequency diversity is easier in WiMAX due to its use of OFDM.

Probably one of the main advantages of WiMAX compared to other 3G technologies is cost. Because WiMAX builds on a lightweight IP infrastructure, the complexity of the core network is reduced. This means that there is no need to maintain both a circuit-switched and a packet-switched core, as is the case in a typical cellular network.

In the end, WiMAX and 3G may become complementary technologies that both address two different sectors of the same market. 3G technologies were developed primarily for voice services while WiMAX was developed with data services in mind. Both are capable of doing both voice and data, but they were not originally designed to do so [35]. This may cause operators to deploy WiMAX as a means to achieve higher data rates while maintaining 3G for voice services.

2.9.2 Wi-Fi

Wi-Fi is based on the IEEE 802.11 family of standards and is primarily a local area network (LAN) technology designed to provide wireless access inside buildings. It operates in licence-exempt bands, and due to interference issues, its allowable transmit power limits are very low. As a result, it only provides indoor coverage of some tens of metres. The current systems based on IEEE 802.11a/g provide a theoretical throughput of 54 Mbps, while only about 20

Mbps is achieved in practice due to overhead from the carrier sense multiple access (CSMA) protocol. The new IEEE 802.11n amendment provides gross throughput of more than 100 Mbps using MIMO technology, and is also expected to provide significant improvements in range through the use of diversity techniques [3].

Even though both the Wi-Fi and WiMAX standards were developed by IEEE, they target different markets. While WiMAX is designed for wireless broadband access, Wi-Fi is not. However, in the past couple of years, many organisations and companies have extended the indoor Wi-Fi coverage to deployments outdoor. Wireless Trondheim is an example of such, where access is provided to the users through dense deployments of base stations in the city. However, because these access points are operating in licence-exempt bands, the allowable transmit power limits the coverage area to some hundreds of meters outdoor. This requires very dense deployment, which is costly and not practical for large scale operation.

Other drawbacks of using Wi-Fi for wireless broadband access is its inefficient CSMA protocol, interference constraints, and its lack of support for high speed mobility[8]. There is also no guaranteed QoS profiles implemented, even though the IEEE 802.11e amendment partly fixes this issue. Wi-Fi was designed to support only a limited number of users simultaneously and does not have the high grade of availability and reliability that is required for carrier-grade technology.

2.9.3 WiBro

WiBro is an acronym for wireless broadband and is the service name for mobile WiMAX in South Korea. It was developed by the South Korean telecommunication industry, which adopted a specific iteration of the IEEE 802.16e-2005 amendment in early 2004. Later that year, it was approved by the Telecommunications Technology Association (TTA) of Korea [12, 33].

Since then, the WiBro community and the WiMAX Forum have decided to co-operate and have been actively working towards aligning the two branches of the standard. It is important to realise that WiBro is based on the same IEEE 802.16 standard as mobile WiMAX, and both the PHY and MAC layer specifications are identical. As of today, WiBro equipment is certified using

[8] The IEEE 802.11p draft amendment is targeting data exchange between high-speed vehicles and between vehicles and roadside infrastructure. This may improve its support for mobility.

a mobile WiMAX certification profiles that specifies a channel bandwidth of 8.75 MHz in the 2.3 GHz frequency band. The fact that WiBro and mobile WiMAX are one and the same technology leads to interoperability and global roaming for users.

2.9.4 802.20

In late 2002, the IEEE 802.20 working group was formed to address a new standard for mobile broadband wireless access (MBWA). It aims to prepare the specification for an efficient packet based air-interface optimised for the transport of IP-based services. The initial design requirements were the following [36]:

- Optimised for vehicular mobility up to 250 kmph
- Operation in licensed frequency bands below 3.5 GHz
- New MAC and PHY specifications without backwards compatibility
- Latency less than 10 ms
- Long range

The purpose and scope of the IEEE 802.20 standard are very ambitious. It aims to fill the gap between cellular networks, which have low bandwidth and high mobility, and IEEE 802 wireless networks, which high bandwidth and low mobility.

The differences between the IEEE 802.20 and IEEE 802.16e-2005 are less pronounced, as both seek to support mobility [35]. However, there are some differences in their approach. IEEE 802.16e-2005 also aims at providing mobility, but this feature was not included from the start as WiMAX was originally meant for fixed access. Mobility was added at a later stage through an amendment. This is not the case with IEEE 802.20, which has been designed for mobility from the ground up. It is being developed specifically for operation in high-speed environments and supports movement of up to 250 kmph. This makes the technology suitable for providing access to trains and other fast moving vehicles. Also, IEEE 802.20 supports ranges that are on the order of cellular systems with very low latency.

Another distinct difference is that IEEE 802.16e-2005 is to a large degree backwards compatible with earlier versions of the standard. IEEE 802.20 is being designed without the constraints of maintaining backwards compatibility, which have several advantages.

In June 2008, the IEEE 802.20 standard was approved. It specifies the PHY and MAC layers of an air-interface that operates in licensed bands below 3.5 GHz. Peak data rates of 1 Mbps per user is supported at vehicular speeds of up to 250 kmph.

2.9.5 802.22

The IEEE 802.22 working group aims at constructing a wireless regional area network (WRAN) that uses cognitive radio to take advantage of white spaces in the allocated TV frequency spectrum [3]. One key feature of the WRAN is its distributed sensing, where connected CPEs periodically report to the BS what they sense in the spectrum. The BS evaluates all the information it receives and decides if a change in frequency must occur.

Because IEEE 802.22 seeks to operate in the very high frequency (VHF) and ultra high frequency (UHF) bands, radio propagation conditions are favourable, providing fairly good coverage over long ranges.

The work group is still at an early stage of development, and has not yet released any drafts or standards.

Chapter 3

Wireless Trondheim

This chapter provides a short introduction to Wireless Trondheim, its history, and the technology on which its network is operating. A few usage scenarios for mobile WiMAX are also presented.

3.1 History

Wireless Trondheim is a joint research and development project initiated by the Norwegian University of Science and Technology in 2005. The project aims to cover the city of Trondheim with outdoor wireless Internet access, accessible to everyone. By doing so, the project seeks to father a world leading facility for research and development in the area of wireless technologies, products, and services.

The idea of establishing a city-wide wireless network in Trondheim initially originated from the research community at NTNU in early 2005. By late November the same year, several other parties also joined the project. During the following months, a small test-bed was set up using equipment from Airespace, and valuable knowledge was gained from this pilot. The project continued to grow, and by September 2006 most of the inner city of Trondheim was covered. At the same time the network opened its services to the students and affiliates of NTNU and the other project members.

As of September 2007, the networks of Wireless Trondheim is available to the general public based on a subscription where the user pays a small fee for twelve hours of Internet access. Students and affiliates of the project members are still able to access the network for free.

3.2 Technology

The networks of Wireless Trondheim have been built using well-known and widespread technologies. The clients connect to the access points using IEEE 802.11a or 802.11b/g, which is a de-facto standard for wireless communication. The access points are either fed with fibre links or connected to the backhaul using high capacity 802.11a radio links.

The platform upon which Wireless Trondheim is built supports technologies that are well suited for testing new and innovative services [37]. The network handles mobility at layer 2, where a user may seamlessly roam from one access point to another without losing connectivity. All access points are connected to the backhaul using at least 10 Mbps symmetric connectivity, facilitating testing of high bandwidth demanding services. Location information is also available as the network gathers information about where a particular user is and offers the possibility of providing location dependent services. The network also supports all well known and widespread standards for security, including WPA/WPA2 and different EAP authentication schemes using the 802.1x standard.

3.3 Infrastructure

Wireless Trondheim uses mostly equipment from Cisco Systems in their networks. The end system consists of Cisco Aironet 1010 and 1030 lightweight access points that are managed by two Cisco 4404 wireless LAN controllers (WLC). The underlying protocol used to facilitate communication between the WLAN controllers and the access points is the lightweight access point protocol (LWAPP). LWAPP is used to control multiple remote access points from one or more controllers. It allows administrators of the network to more easily configure, monitor and troubleshoot the wireless infrastructure through a centralised interface. Client WLAN traffic is also tunnelled to and from the centralised WLCs using LWAPP.

The network supports several SSIDs, allowing different ISPs to offer their services in the same network. This means that the cost of operating the physical network can be shared among many service providers.

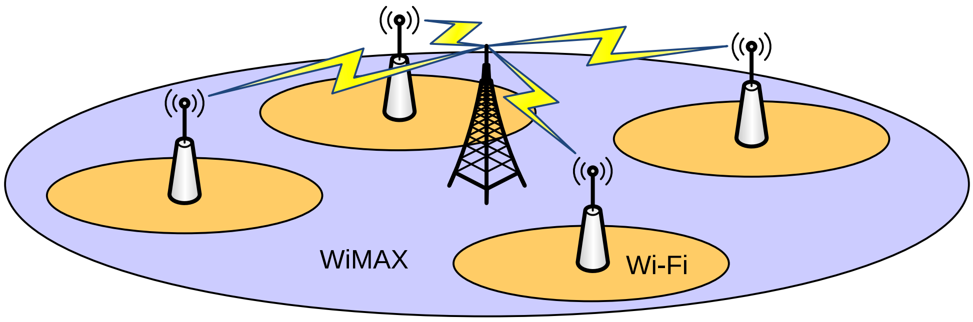

Figure 3.1: Feeding Wi-Fi access points

3.4 Mobile WiMAX Scenarios

A mobile WiMAX network in Trondheim would enable quite a few interesting areas of usage. In this section, three possible scenarios are outlined.

3.4.1 Feeding Wi-Fi Access Points

One possible scenario is to provide backhaul for Wi-Fi access points operated by Wireless Trondheim. As of today, mainly two methods are used for connecting the access points to the core network. They are either connected directly using fibre, or they operate in mesh mode using IEEE 802.11a as backhaul. However, it is neither economically nor practically feasible to deploy fibre to each access point, and relying on other access technologies such as DSL may not be optimal with regards to speed, cost, and ownership of infrastructure. Using IEEE 802.11a for backhauling is not optimal either. First, even though most clients use IEEE 802.11b/g to connect to the network today, IEEE 802.11a may be utilised more in the future. This will lead to shortage of spectrum and increased interference. Second, because unlicensed spectrum is used, it is easier to jam and snoop on the wireless traffic. Third, the output power is limited, hence the mesh links cannot travel over too long distances. And fourth, the CSMA protocol is still used for obtaining access to the medium, which implies that it is difficult to implement strict QoS on different traffic streams traversing the wireless link.

Instead, WiMAX may provide a more optimal solution for backhauling, as illustrated in figure 3.1. This configuration is particularly suited for areas where the aggregate load on the Wi-Fi access points is relatively small. Because WiMAX is capable of covering a much larger cell size than Wi-Fi, it

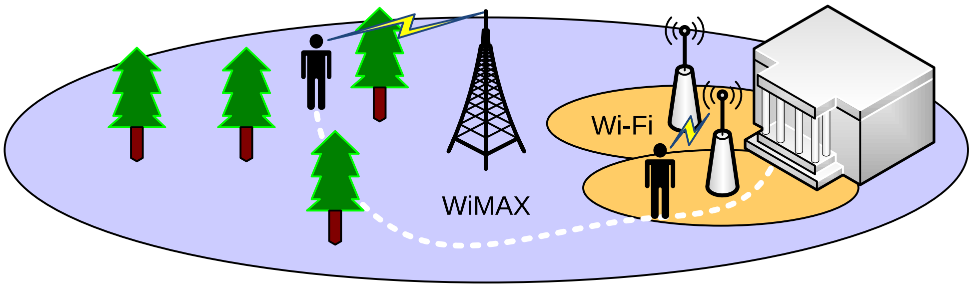

Figure 3.2: Providing access outside Wi-Fi coverage

may actually provide backhauling for multiple Wi-Fi access points simultaneously. The advantages to WiMAX versus Wi-Fi were discussed in section 2.9.2 and are not repeated here. But most importantly are the in-built support for strict QoS, operation in licensed bands, longer reach, and higher grade of reliability.

3.4.2 Providing Access outside Wi-Fi Coverage

Wi-Fi access points cover an area of limited size, mostly due to their requirements for high throughput and limited equivalent isotropically radiated power (EIRP). While it is feasible to cover indoor areas and dense outdoor areas using Wi-Fi, it soon becomes costly and impractical to cover large outdoor areas. In sub-urban areas, parks, or sparsely populated areas, it may be more practical to provide umbrella coverage using WiMAX.

The idea is that a user primarily connects to a Wi-Fi access point if there is coverage, and then uses inter-media roaming to WiMAX when moving outside the Wi-Fi coverage area. This would require the subscriber unit to support both Wi-Fi and WiMAX, but such functionality is on the horizon. This way, the user would stay connected, albeit the throughput of the WiMAX network would not be able to compete with the throughput of the Wi-Fi network. This scenario is illustrated in figure 3.2.

3.4.3 Providing Access on Public Transportation

The third scenario is based on the idea of providing access to public transportation, such as buses, trams, and ferries. This would be quite difficult to achieve using Wi-Fi only, as the number of access points needed to cover the service area would be prohibitively large. More so, Wi-Fi is not designed

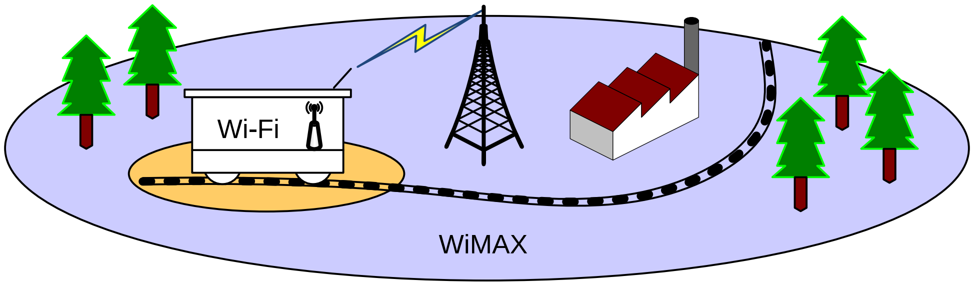

Figure 3.3: Providing access on public transportation

for neither backhauling nor mobility, whereas WiMAX is. This scenario is illustrated in figure 3.3. Here, a local Wi-Fi access point is operating within the tram, providing access to the passengers. The access point is fed using WiMAX, where the base station is connected to the core network via a high capacity fibre link.

The main focus on this thesis is to investigate the possibility of using mobile WiMAX to provide Internet access on the tram in Trondheim. In the next chapter, the process of planning such a network is outlined.

Chapter 4

Network Planning

This chapter deals with the network planning stage. First, the physical area to cover is identified and potential base station sites are located. Next, the knowledge gained from site surveys at two possible locations are outlined. Finally, the results from propagation modelling are presented along with recommendations for improving the coverage.

4.1 Site Planning

The first stage in the process of network planning is to identify the physical area to cover. As there is only a single tram line in Trondheim, this is relatively straight forward. The line, which is depicted in figure 4.1, starts in the city centre and runs to Lian. The entire journey is 9 km and takes about 20 minutes from start to end.

Next, suitable locations for base stations must be identified. There are a few requirements that must be met for the ideal site. First and foremost, the site must be available for lease and the least cost must be acceptable. Second, the site must be elevated above surrounding buildings and provide adequate coverage for the intended area. Third, installation of the base station equipment and antennas must be feasible. This includes secure and easy access to the roof with enough space for the equipment. Fourth, there must be easy access for power and backhaul such as fibre [21]. As the cost of establishing a site is high, the number of base stations should be kept as low as possible, while still maintaining full coverage of the tram line.

In figure 4.1, three potential sites have been identified. The first site is

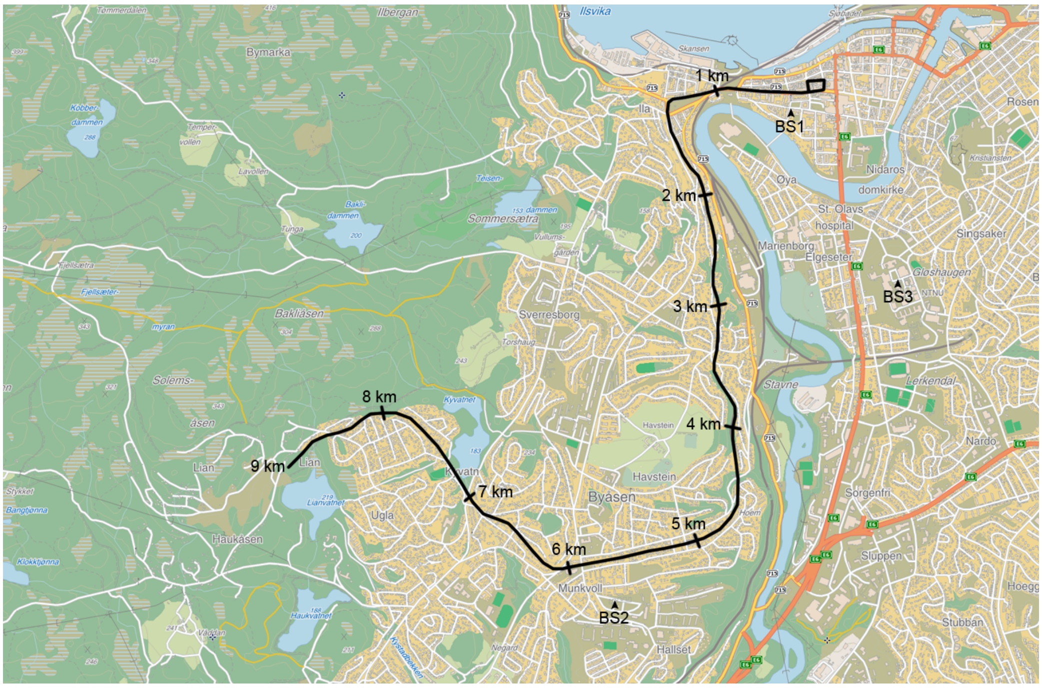

Figure 4.1: Map of the tram line [38]

located at the Gunnerus Library in the city centre, a building housing the oldest scientific library in Norway. It is denoted BS1 in the figure. Because the premises are owned by the university, site access is easy and lease cost is affordable. There is also adequate space on the roof for installing the equipment, and fibre backhaul is readily available.

The second site is located at Byåsen upper secondary school, denoted BS2 in the figure. The premises are owned by the county council, which is also co-owner of Wireless Trondheim. Hence, obtaining access to the roof of the building would not be too problematic. There is power readily available, and fibre infrastructure could easily be extended to the rooftop.

The third potential site is located at the university campus. Site lease would not be an issue, and both power and fibre infrastructure would be easily available. From this location, there is a clear line of sight to almost half the tram line. However, the distance from the site to the line is at least 1.5 km. This is a main drawback, as much of the emitted energy from the antennas would cover an area with limited potential users.

Figure 4.2: View westwards from BS1

4.2 Site Survey

Three potential sites have been identified. The next step is to perform physical site surveys to get a visual understanding of the area to cover. For the purpose of this project, it is important to note high buildings that block direct line of sight to the tram line as well as hills and dense trees that could impact signal propagation. The building roof must also be inspected to find suitable places to put the antennas and the base station equipment. Cable ducts must also be located if power and fibre connections are to be deployed.

Two site surveys were carried out, and they are detailed in the following sub-sections.

4.2.1 Site BS1

The library is considerably higher than most of its surrounding buildings. An antenna mounted here could in theory provide coverage to the entire inner city of Trondheim. The rooftop has easy access and there is plenty of space both for the base station equipment and the antennas. Power is already available, and it is relatively straight forward to install fibre backhaul as well.

With regards to the tram line, some parts of the track are clearly visible from the base station while others are not. This is illustrated in figure 4.2, which is a photograph taken due west from the site. As can be seen, the solid line shows where the line runs in the open, while the dotted line shows where the line has no line of sight to the base station.

The tram line starts to the right of the picture, and for the first kilometre it passes through non line of sight areas. The distance from the base station to the line varies between a few hundred meters up to a kilometre at the point where the line enters into the open area. In order to achieve good coverage

Figure 4.3: View north-eastwards from BS2

on this first stage, it is vital that the signals diffract over the roof edges of the nearby buildings and into the streets below. Because the density of the buildings is not too high in this part of the town, signals may also propagate by reflection.

When the tram line enters into the open area, it continues to be more or less exposed for line of sight communication for the next three kilometres. This can be seen in figure 4.1 and 4.2, where the line continues southwards along the hill. The distance from the base station varies between one and three kilometres, and there are no dominant obstacles in the path. However, vegetation along the track is not negligible as some parts are covered with dense trees. This is not easily seen in the photograph, as it was taken during early spring.

The conclusion from the site survey is that it should be possible to cover the first part of the track from this site using a single 120 degrees sector focusing south-westwards. This is illustrated in figure 4.4. The antenna could optionally be placed below the roof to prevent interfering with potentially other sectors east of the site. A down-tilt of a few degrees should be applied to make sure its main lobe hits the track and not the hills above. If the antenna is mounted below the roof, its back lobe would probably not provide adequate coverage along the first hundred meters of the track. In this case, another sector pointing eastwards may be necessary to provide full coverage.

4.2.2 Site BS2

The school is located on a relatively flat area and is taller or at least as tall as its surrounding buildings. The rooftop of the building is easily accessible and there is enough space for placement of the base station equipment. Power is also readily available and there should be no problems deploying fibre through existing cable ducts.

From the top there is a clear view in all directions, and the surrounding area is considered as sub-urban with a medium density of trees and other vegetation. A photograph taken north-eastwards is shown in figure 4.3. Here, the tram line is marked as a black line passing from right to left. As can be seen, the tracks run relatively close to the premises, with a minimum distance of about 350 meters. A few medium sized buildings are located between the site and the tram line, hence non line of sight communication is necessary. However, because the area is very open, the radio signals should propagate well via multipath and reflection and still provide good coverage along this part of the track. The end station is located behind the hill spotted to the left of the picture.

There might be some problems achieving good signal coverage to the last part of the tram line. Especially the last kilometre may cause some challenges. Here, the tram passes behind a hill known as Ugla, and the area is covered with high density of trees and buildings. The distance between the base station and the track here is relatively long, varying from 1 to 2.5 kilometres.

It might be possible to cover the remaining parts of the tram line using a single sector of 120 degrees pointing northwards. However, this is not an optimal solution for at least two reasons. First, the main lobe of the antenna would only cover a small area of the track. Second, because the site is located atop of a hill, directing the antenna beam north- or eastwards would cause much of the signal energy to hit the city below, and thus cause interference with other potential cells using the same frequency.

Instead, the sector should point north-westwards to focus its main energy towards the problematic area mentioned above. The antenna should be mounted on a pole to gain a few meters extra height, and its main lobe should be down-tilted by a few degrees. This solutions leads to the requirement of a second sector pointing north-eastwards to provide adequate cover to the remaining part of the track. As mentioned above, one should be careful to avoid interfering with other potential cells below, thus reduced transmit power and severe down-tilt may be required.

A sketch of the planned sectors from both sites are outlined in figure 4.4.

4.3 Propagation Modelling

When designing a wireless network, it is imperative to gain thorough understanding of the coverage provided by each site. In order to evaluate and

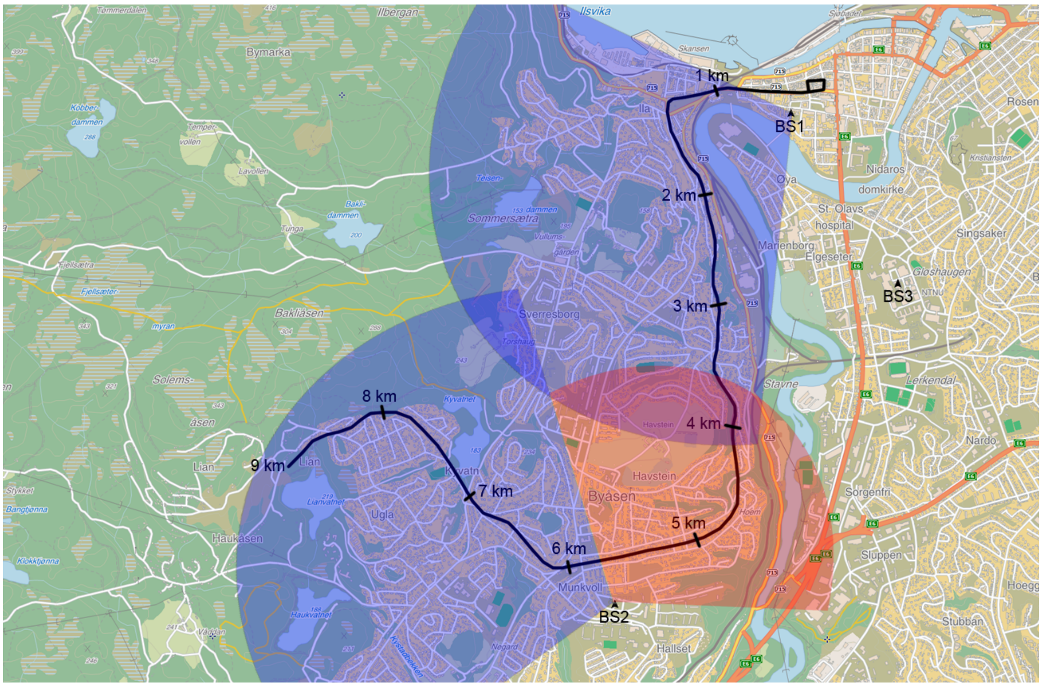

Figure 4.4: Map of the tram line with planned sectors. Note that the sector north-eastwards from BS2 is smaller than the others to indicate reduced transmit power

estimate the coverage, either propagation modelling or physical surveys may be used. Propagation modelling is accomplished with a software tool while physical surveys are accomplished by temporarily installing the base station equipment and then measuring the resulting coverage. Even though physical surveys are more accurate, in the case of a large network, the time and expense associated with physical surveys at multiple sites may not be worthwhile [21].

The propagation modelling software may use a variety of radio propagation models to estimate the coverage area. The following sub-section provides a brief introduction to these models.

4.3.1 Radio Propagation Models

During the 70s, a significant body of work defining the properties of RF propagation was accomplished. This work led to the development of a series of algorithms that describe the mean behaviour of RF propagation over varying environments and terrain. These empirical propagation models are, at best, estimations of real-world propagation as they are based upon statistical behaviour. A number of different models exists, and the following list shows commonly used models for outdoor environments:

- ITU-R
- Longley-Rice model
- Okumura model
- Okumura-Hata model for urban, sub-urban, or open areas
- Cost-231 model

Next, some of these models are outlined in more detail.

Okumura-Hata

The Okumura-Hata models extend the Okumura model to realise the effects of diffraction, reflection, and scattering caused by city structures. Three different models have been developed for different environments: urban areas, sub-urban areas, and open areas [3]. The model for urban areas is formulated as follows:

$$OHUA_{dB} = 69.55 + 26.16 \log f - 13.82 \log h_B - C_H \\ + [44.9 - 6.55 \log h_B] \log d \tag{4.1}$$

where f is the frequency in MHz, h_B is the height of the base station antenna in metres, and d is the distance between the base station and mobile station in kilometres. C_H is a correction factor for the antenna height of the mobile station, and is defined as follows for small or medium sized cities:

$$C_H = 0.8 + (1.1 \log f - 0.7)\, h_M - 1.56 \log f \tag{4.2}$$

where h_M is the height of the mobile station antenna in metres. There also exists another correction factor for larger cities.

The Okumura-Hata model for sub-urban areas is given by:

$$OHSA_{dB} = OHUA_{dB} - 2\left(\log \frac{f}{28}\right)^2 - 5.4 \tag{4.3}$$

The Okumura-Hata model for open areas is given by:

$$OHOA_{dB} = OHUA_{dB} - 4.78\,(\log f)^2 + 18.33 \log f - 40.94 \tag{4.4}$$

Cost-231

The Cost-231 model extends the Okumura-Hata models to cover higher frequencies. It is formulated as follows:

$$COST231_{dB} = 46.3 + 33.9 \log f - 13.82 \log h_B - C_H \\ + [44.9 - 6.55 \log h_B] \log d + C \tag{4.5}$$

where f is the frequency in MHz, h_B is the height of the base station antenna in meters, and d is the distance between the base station and mobile station in kilometres. C_H is identical to the correction factor used in the Okumura-Hata model for urban areas and C is a constant equal to 0 dB for medium sized cities and sub-urban areas, and 3 dB for metropolitan areas.

Free Space Path Loss

The free space path loss (FSPL) model is not a radio propagation model per se, but it is included here because it is commonly used to predict the loss of signal strength for a radio wave propagating through free space. The model states that the power loss is proportional to the square of the distance between the transmitter and receiver. During the testing phase, the model was used to compare the received signal strength to the distance between the client and the base station, hence it is introduced in this section. The model, in terms of dB, is formulated as follows:

$$FSPL_{dB} = 20\log(f) + 20\log(d) + 32.44 \tag{4.6}$$

where f is the frequency in MHz and d is the distance in km.

Comparison of the Models

The power loss of the four aforementioned models as a function of distance are graphically illustrated in figure 4.5. The topmost line illustrates the loss according to the COST-231 model from equation (4.5) with C equal to 0, while next upper line illustrates the loss according to the Okumura-Hata model for urban areas from equation (4.1). The middle line represents the Okumura-Hata model for sub-urban areas from equation (4.3), while the next lower line shows the Okumura-Hata model for open areas from equation (4.4). Last, the bottom line represents the free space path loss from equation (4.6). For all models, the following parameters have been used: Frequency f equal to 2680 MHz, H_b equal to 20 metre, and H_m equal to 2 metre.

It is clearly evident that the losses calculated by the COST-231 model and the Okumura-Hata model for urban areas coincides to a large degree, with a difference of about 3 dB. These models also predict the highest loss, which makes sense as they are built for operation within dense cities.

4.3.2 Cartographic Information

In their simplest form, the propagation models can be used as is to calculate a rough estimate of the coverage. This is useful for planning the number of sites necessary to cover a specified area. However, in order to provide accurate results, more complex models are necessary that utilise knowledge of

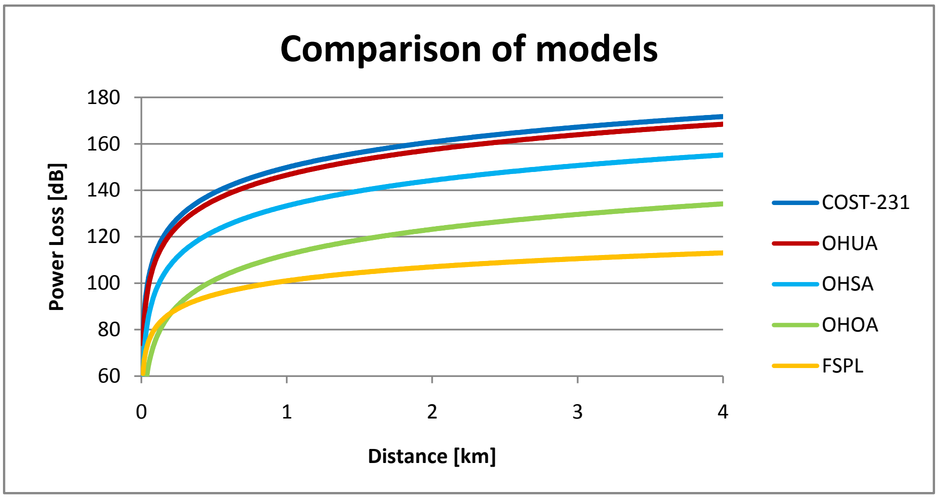

Figure 4.5: Comparison of path loss models

the terrain and ground occupancy. Commonly, three different types of cartographic layers are used for propagation modelling: a digital terrain model, a clutter layer, and a building height file.

The resolution of the layers specifies the level of detail that is available. For instance, a 25 meter digital terrain model averages the terrain in a 25 x 25 meter square which is represented as a single elevation. Naturally, this averaging leads to some inaccuracy in coverage prediction. There are digital terrain models available for all areas of the planet. Some are free of charge, while the more detailed ones come with a high price tag.

The digital terrain model only specifies the elevation of the terrain, not what occupies the ground. The clutter layer, or ground occupancy layer, defines the morphological data of the terrain such as forestry, parks, roads, buildings, water, and open areas. The clutter layer is also provided with a specific resolution, where it is imperative to obtain as detailed data as possible for the propagation modelling to be realistic. In addition to the clutter layer, a building height file may be used to provide even more accurate results. In this map layer, each building is modelled with its exact height, providing a detailed environment. Morphological data is somewhat time limited in its utility as areas change. Hence it is necessary to ensure that both the clutter layer and building height files are up to date.

4.3.3 Modelling Software

There are a number of different software packages that are suitable for propagation modelling. For the purpose of this project, ICS telecom nG from ATDI has been used. Another option is Radio Mobile, which is a freeware tool.

Note that acquiring a software package and operating it is not the hard part, effectively using it is. There are a number of issues to consider in order to assure that the modelled propagation reflects the real world. First, it is vital to obtain high resolution terrain and morphological data that is accurate and up to date. Next, it is important to make sure that all the parameters that are put into the modelling algorithms are as close to reality as possible. It is common procedure to perform field measurements using real equipment in parallel with the propagation modelling, and compare the results. This comparison allows tweaking of the parameters in the model for the prediction to line up with reality [21].

ICS telecom nG

ICS telecom nG [39] is a tool for planning telecommunication networks and managing frequency spectrum. It handles frequency ranges from 10 kHz up to 450 GHz and provides an extensive set of features. For instance, it may be used for frequency planning and spectrum optimisation, network design, coverage analysis, and interference calculation. The software supports five cartographic layers:

- Digital terrain model
- Geo-referenced maps or imagery
- Clutter files
- Vector files
- Building height files

The propagation modelling has been carried out in version 8.8.7 release 1258 as of November 14 2007.

Radio Mobile

Radio Mobile [40] is a widely used freeware modelling program developed by Roger Coudé. It is available for the Windows operating system and new versions of the software that corrects bugs and adds new features are continuously made available. The latest version is 9.1.9 as of June 17, 2008.

In addition to the modelling software, elevation data is necessary. Such data covering the whole earth is available free of charge from different sources, among others the Shuttle Radar Topography Mission (SRTM) web page [41]. The resolution of the free data typically resides in the area of 3 arc seconds (100 meter) and upwards. Unfortunately, terrain data for the area of Trondheim is not available from the SRTM as it only covers up to about 60 degrees north. Trondheim is located at 63 degrees north.

4.3.4 Modelling Setup

Propagation modelling is accomplished by first configuring some general parameters for the scenario. Next, the base stations must be configured and placed at their respective sites. Finally, the customer premises equipment (CPE) must be configured. In the following, each of the steps performed in ICS telecom nG are outlined in more detail.

General Parameters

The first thing that needs to be done after starting the program is to load the different mapping layers. In this project, both a 25 meter digital terrain model and a 100 meter clutter layer were available for the desired area. A building height layer was not used, hence signal propagation was modelled only based on the elevation of the terrain and the clutter layer which describes the ground occupancy.

Preferably, the Cost-231 model should have been used for propagation prediction. Unfortunately, the model did not produce realistic results, and because of time limits the Okumura-Hata model for sub-urban areas was chosen instead. It is commonly used for predicting the behaviour of cellular transmissions inside small cities. Even though the model is designed for frequencies up to 1.5 GHz, it may be used for higher frequencies if correction factors are included.

Parameter	Value
Nominal power	2.5 W
Antenna gain	17 dBi
System loss	1 dB
EIRP	99.5 W
Frequency	2680 MHz
3dB azimuth beamwidth	60 degrees
Antenna height	20 m
Bandwidth	5000 kHz

Table 4.1: Base station parameters

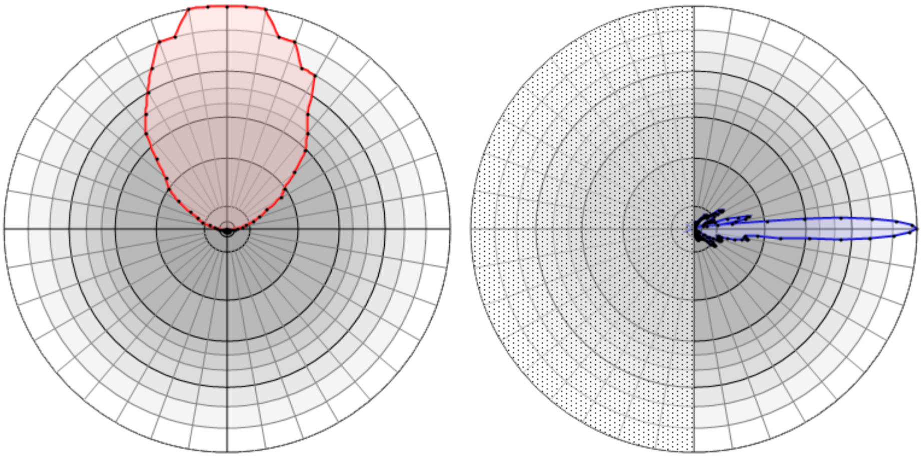

Figure 4.6: Base station antenna pattern

Base Station Parameters

The base stations were configured with the parameters shown in table 4.1. The output power of the transmitter was set equal to 2.5 W, or 34 dBm. With an antenna gain of 17 dBi and 1 dB loss, the total EIRP used in the modelling was 50 dBm. This is much less than the limit in the spectrum licence, which is equal to 54 dBm/1MHz. A single 60 degree antenna was used per sector, and no diversity gain nor down-tilt were configured. The antenna radiation pattern is illustrated in figure 4.6. It is more or less identical to the actual antenna that was used in the field test, which makes the modelling more realistic.

Modelling was carried out with base stations located at each of the three potential sites that were identified during the initial site planning stage.

Parameter	Value
Antenna gain	6 dBi
System loss	1 dB
Frequency	2680 MHz
Antenna height	2 m
Sensitivity	-77, -86, -97 dBm

Table 4.2: CPE parameters

CPE Parameters

The CPE was configured with the parameters as shown in table 4.2. The sensitivity levels are according to the product specification for the actual CPE that was used in the field test. It specifies that typical sensitivity values are -77 dBm for 64-QAM 3/4 and -97 dBm for BPSK 1/2, both at 5 MHz bandwidth with a BER of 10^{-6}. This translates to an achievable gross throughput of 18 Mbps for the highest and 2 Mbps for the lowest modulation. Because the CPE should be mounted on a tram, it was decided that an omni-directional antenna had to be used. A typical value for such antenna is 6 dBi.

4.3.5 Modelling Results

A number of different modelling scenarios were carried out for each of the potential sites. Some of the more interesting ones are presented in the following, while more are available in the appendices.

Site BS1

The first potential site is located atop of the library in the inner city. It was set up with one 60 degree sector pointing south-westwards with an azimuth of 240 degrees and no down-tilt. The modelling result is shown in figure 4.7 and is interpreted as follows. The different shades of colours illustrate the varying signal strength. For instance, the blue areas represent a signal strength of -96 dBm, and is close to the weakest signal the CPE may operate under. In other words, in these areas expected gross throughput is 2 Mbps or less. In areas that are green or warmer, the signal strength is higher and more throughput may be expected.

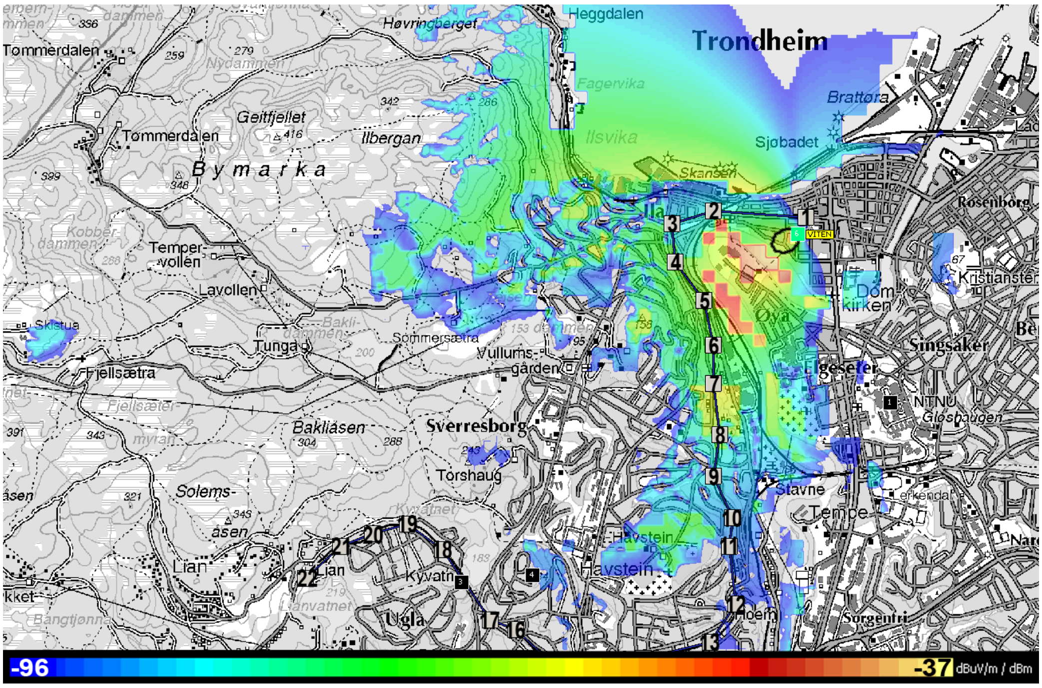

Figure 4.7: Estimated coverage from BS1 with -96 dBm sensitivity

The current configuration estimates that approximately half the track can be covered from this site. However, there are some exceptions that need to be clarified. First, the start of the tram line, which is marked as point 1, lies outside of the expected coverage area. Second, problems related to weak signal reception may arise in the area between point 1 and 3, and onwards from point 8. Here, the expected signal strength is not adequate to make use of higher orders of modulation, and throughput is expected to be low. From point 12 and beyond the signal strength quickly drops below the noise floor, and communication from here seems unlikely.

Another noticeable point is that the tram runs in non line of sight areas between point 1 and 4. Signal reception here may be much less than predicted by the modelling software as no building height file was used.

Adjusting the azimuth of the antenna does not resolve these problems. Turning it further north causes less signal energy to reach the southern parts of the track, and much energy is spread out to sea. Turning it further south causes much of the energy to interfere with other potential sectors to the south, and causes less coverage at the start of the tram line. The conclusion is that an azimuth of 240 degrees seems to provide best coverage using a single sector. In order to provide better coverage along the first part of the

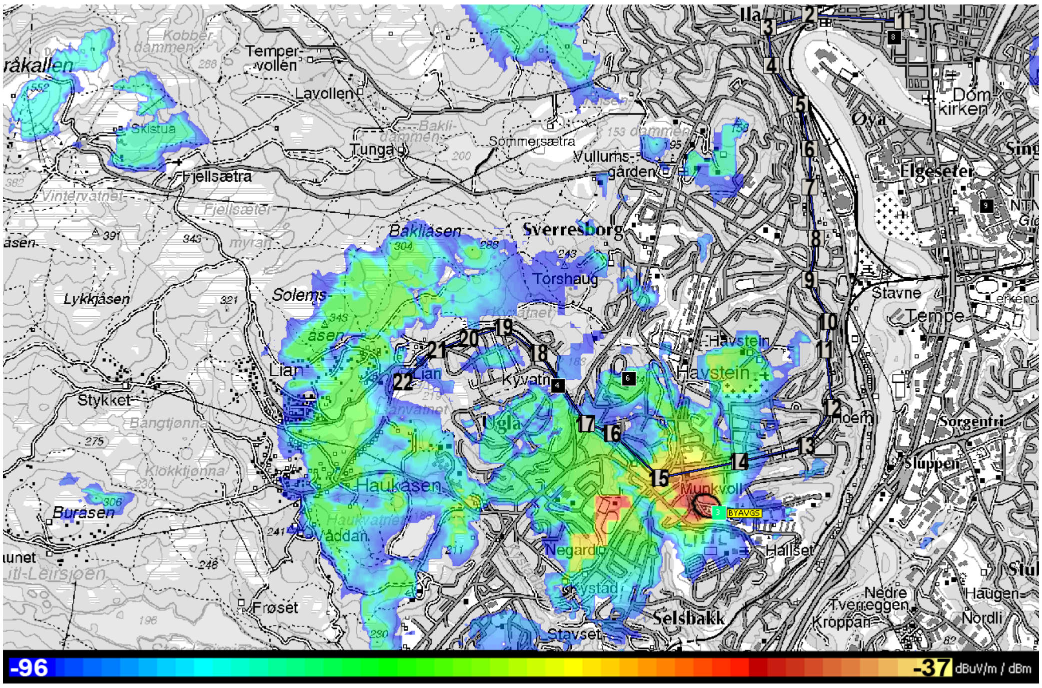

Figure 4.8: Estimated coverage from BS2 with -96 dBm sensitivity

track, it is recommended to make use of a second cell focusing eastwards.

Site BS2

The second potential site is located atop the upper secondary school. This sector is also covering 60 degrees, pointing north-westwards with an azimuth of 300 degrees and no down-tilt. The modelling result is presented in figure 4.8. Here, the signal strength is quite strong between point 14 and 15. In this area it is estimated that the highest order of modulation may be utilised, and hence high throughput may be achieved.

However, there are two problematic areas. The first is the area leading up to point 14, where the signal is too weak for any communication, even using the lowest order of modulation. The reason for this is that the antenna lobe is pointing too much westwards, and little or no energy hits this area. As discussed during the site survey, it seems necessary to add another cell to this site in order to provide adequate coverage.

The other area is the last part of the track, from point 17 and onwards. Here, the modelling results estimate that virtually no communication is possible for the remaining length of the track. The only exception being the end station, where the lowest possible modulation may be used. This area was also identified as problematic during the physical site survey, as the hill known as Ugla blocks signal propagation to the tram line which runs behind it. While the base station is located at 130 meters above sea level, the hill rises more than 220 meters high, effectively blocking propagation to the tram line behind.

It is also worth noticing that there is a small hill right north of point 15 which partly blocks signal propagation to the track behind it. This area is identified as the blue area between point 15 and 17. Along this part of the tram line, only the lowest order of modulation is possible.

Combining BS1 and BS2

In order to gain a better understanding of the composite coverage when using both sectors simultaneously, both BS1 and BS2 have been combined into one figure. The result is illustrated in figure 4.9 and 4.10, where the latter provides a 3D illustration. As can be seen, the two sectors provide pretty good coverage of the track as a whole, with the exceptions that already have been mentioned. Interference measurements were also carried out, with the conclusion that the two sectors do not interfere with one another. Hence, they may use the same frequency.

Note that all the propagation results presented so far use sensitivity values ranging from -96 dBm and upwards. This means that the blue colour represent areas where the signal strength is just adequate for the CPE to connect to the base station and communicate using the lowest order of modulation. However, it is also interesting to identify the areas where the signal strength is strong enough to enable full throughput.

Figure 4.11 represents the composite coverage of the two sectors where the sensitivity value ranges from -76 dBm and upwards. In other words, all colours now represent areas where the signal strength is adequate for the use of the highest order of modulation.

It can easily be seen from the figure that full throughput is only possible between point 4 and 8, and between point 14 and 15. The remaining length of the track lies outside of the coloured area, hence lower throughput may be expected here. It is very interesting to note that full throughput is not possible between point 15 and 17, even though the area is located right in

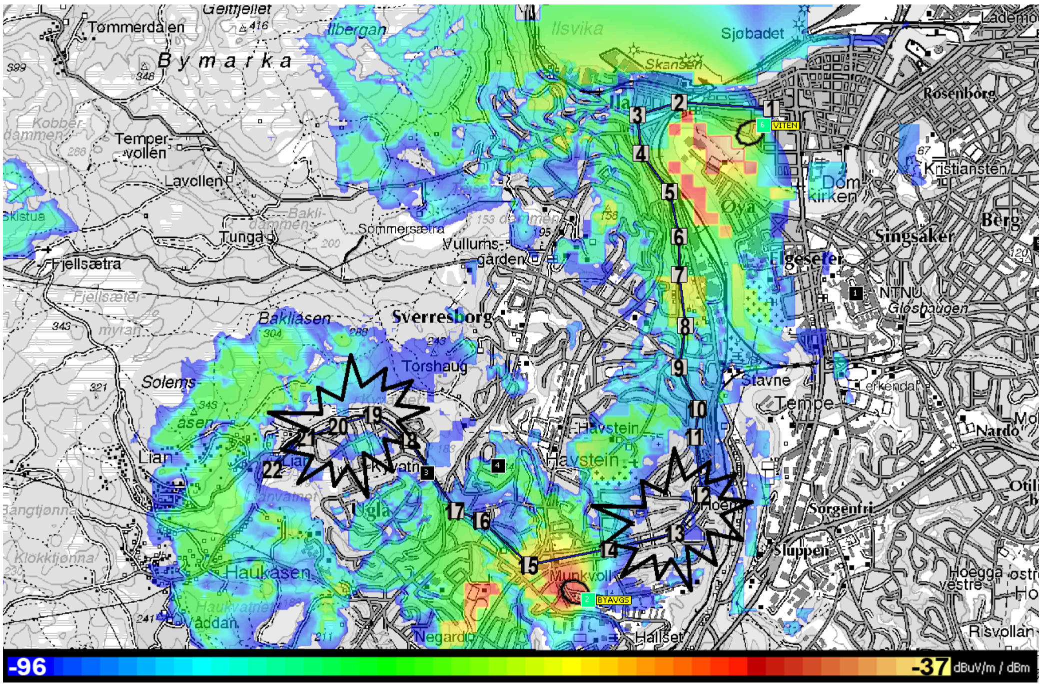

Figure 4.9: Estimated coverage from BS1 and BS2 with -96 dBm sensitivity

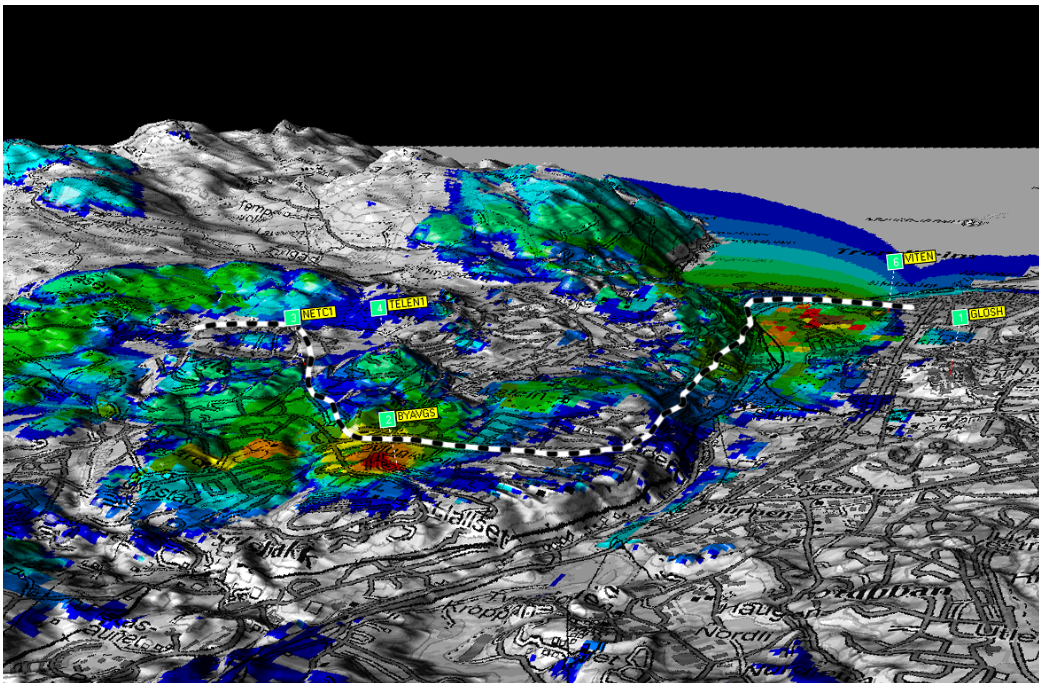

Figure 4.10: Estimated coverage in 3D from BS1 and BS2

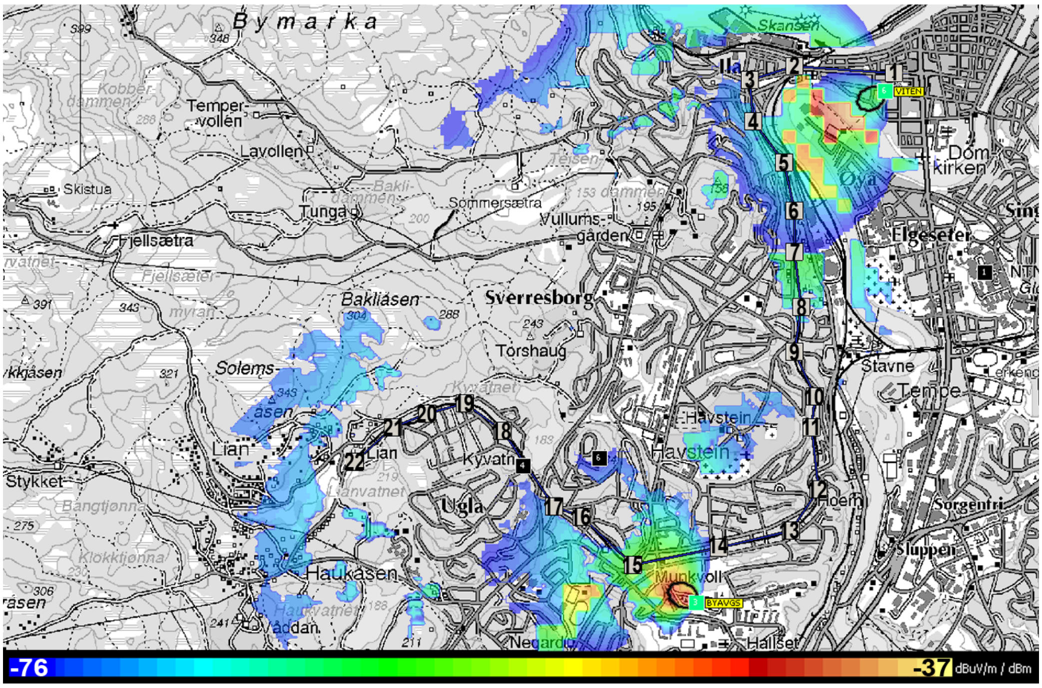

Figure 4.11: Estimated coverage from BS1 and BS2 with -76 dBm sensitivity

the middle of the main lobe. The distance from the base station to this area varies from 0.5 to 1.5 kilometre.

4.3.6 Recommendations for Improvement

4.3.6 Some adjustments need to be made to ensure that the complete tram line is covered. In the following, some recommendations are presented.

Adding a Sector from Site BS3

The initial site planning section identified three potential sites. Thus far, only BS1 and BS2 have been dealt with. In order to provide better coverage between point 8 and 14, a new sector could be installed at site BS3. The potential sector is covering 60 degrees with an azimuth of 235 degrees, thus pointing south-westwards. As can be seen from figure 4.12, it provides coverage along the whole north-south leg of the line. However, it is only between point 6 and 12 that optimal signal strength is obtained.

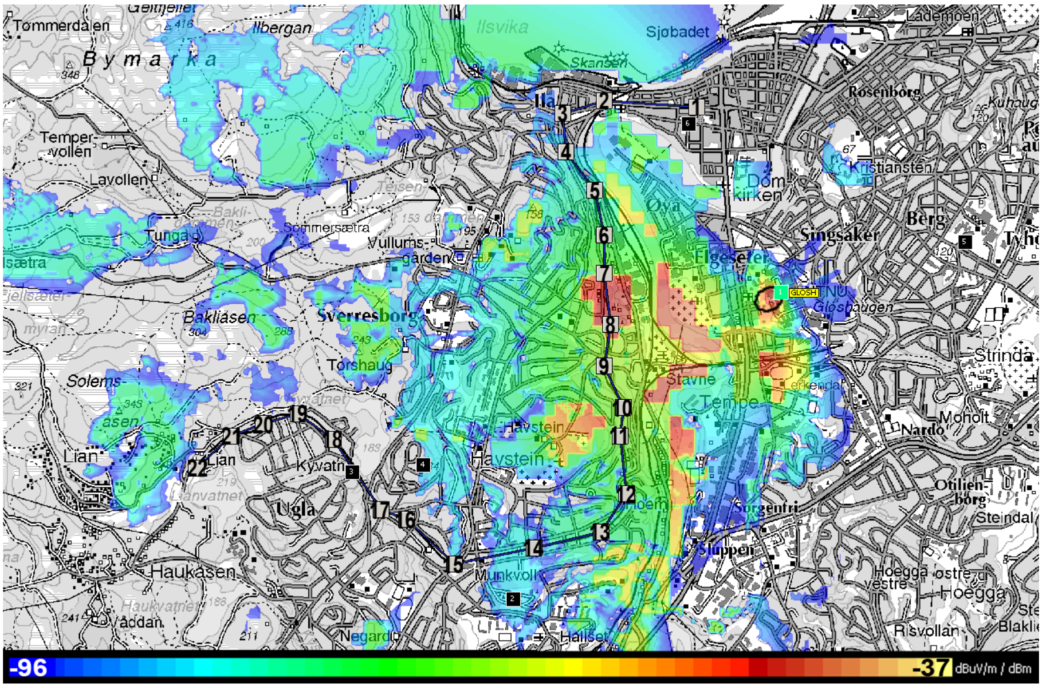

Figure 4.12: Estimated coverage from BS3 with -96 dBm sensitivity

This solution has a few drawbacks. First, it is economically expensive to deploy three sites to cover the tram line. Second, the distance between site BS3 and the line is 1.5 km at its minimum, which is not optimal with regards to path loss. Third, this sector overlaps with the sector from site BS1 to a large degree, and providing double coverage is a waste of resources. Fourth, if the sectors from BS2 and BS3 are operating on the same frequency, there could arise problems with interference.

Adjusting the Sector at Site BS2

Another solution would be to adjust the sector at site BS2 northwards, thus providing better coverage in the problematic area to the east. As shown in figure 4.13, this would cause much better signals between points 11 and 14. However, much of the signal energy would also spread out over the entire city below, which is not desirable.

The conclusion is that the sector at site BS2 cannot be turned much in order to avoid interference problems in adjacent sectors. It might be feasible to implement two sectors at the site, where the second sector provides coverage north-eastwards. In that case it is important that this sector transmits at a

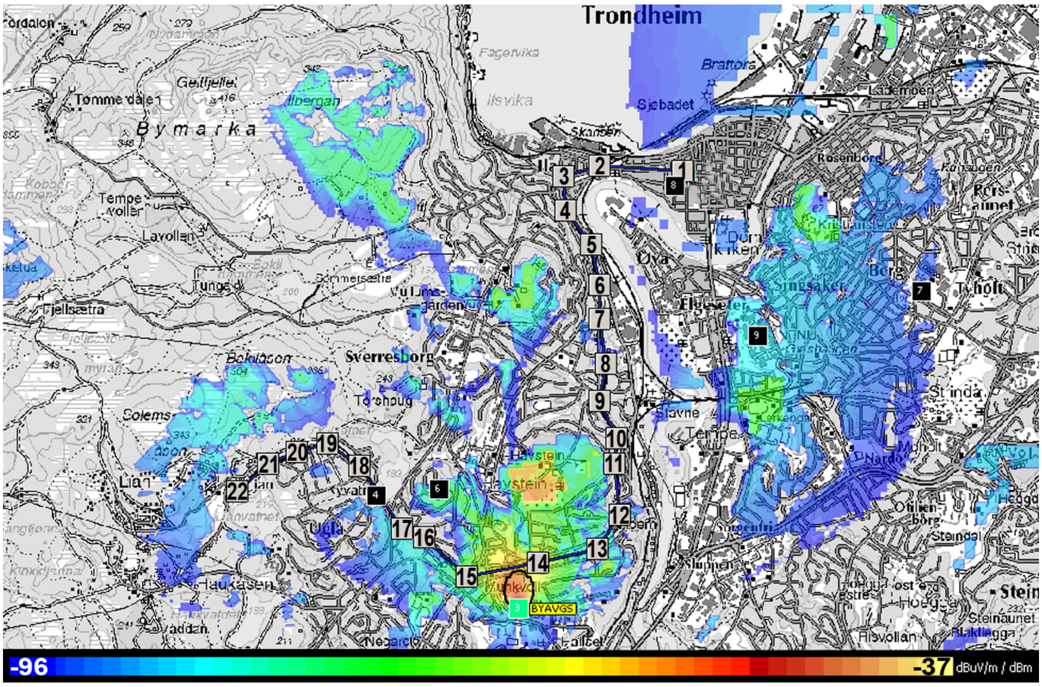

Figure 4.13: Estimated coverage from BS2 with antenna pointing northwards and -96 dBm sensitivity

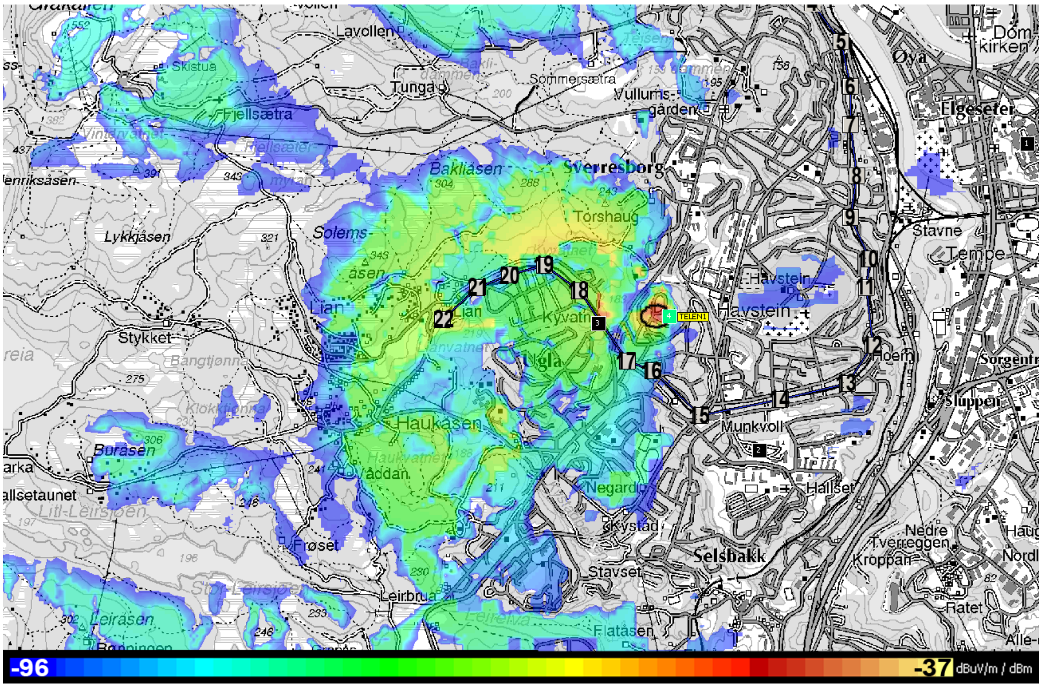

Figure 4.14: Estimated coverage from Storhaug with -96 dBm sensitivity

very low power level to avoid interference in other cells. This scenario has not been modelled.

Adding a Sector from Storhaug

A third option is to establish a new sector from the cellular tower owned by Telenor at Storhaug. This concept is outlined in figure 4.14, where a 60 degree sector is pointing almost due west. From this site, the last part of the tram line is covered easily.

Establishing a new site at this point is not straightforward. First, a lease contract must be arranged with the owner of the tower. Second, depending on the backhaul available at the site, it may be necessary to deploy a radio link feeding the base station. This will increase the cost and complexity, and may not be feasible. Third, the last part of the track is not too important with regards to optimal coverage, as the number of passengers travelling here is very low.

Chapter 5

Network Implementation

This chapter details the network implementation phase of the project. It starts with a short summary of the vendor selection process. Then, the chosen WiMAX system is described in some detail along with specifications for both base stations, antennas, and subscriber units. The process of installing and configuring the system is also included.

5.1 Vendor Selection

A complete WiMAX network solution consists of equipment from different vendors. For instance, some supply only core network components such as switches, routers, or firewalls. Others develop applications and solutions for billing, customer relationship management, and network operations and maintenance. Likewise, some vendors manufacture only base stations while others provide only the subscriber units. Chip manufacturers and antenna makers also provide distinct parts of what is known as the WiMAX ecosystem.

The process of selecting a vendor is not detailed here. However, a short summary is given. At the time of writing the thesis, no products were certified against the mobile WiMAX system profile[1]. There only existed products that were pre-certified, which translates to systems that implement all or parts of the functionality as defined by mobile WiMAX system profile.

[1] In April 2008, the WiMAX Forum announced that the first products had been certified against the 2.3 GHz certification profile. Two months later, in mid June 2008, another ten products were announced certified against the 2.5 GHz certification profile.

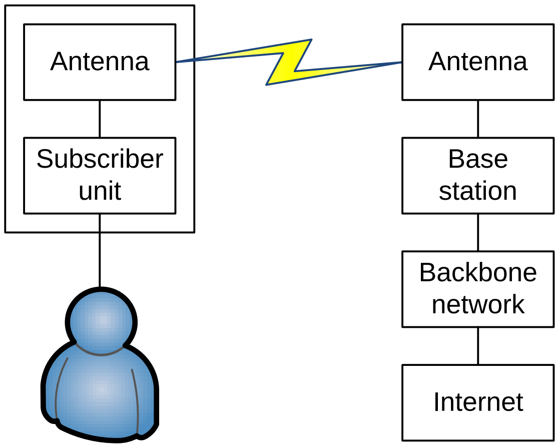

Figure 5.1: High level system description

Equipment from different vendors were evaluated with respect to cost, availability, and functionality. Based on a holistic assessment, Alvarion Ltd. was chosen as supplier for base stations, antennas, and subscriber units. Core network components were delivered by Cisco as part of the ongoing contract with Wireless Trondheim. Unfortunately, the equipment was not delivered before end of May. This led to stringent time constraints with regards to installing, configuring, and testing the equipment before the deadline of submitting this report.

In addition to the physical equipment, a frequency licence was also needed. The process of obtaining a licence is not detailed here, but an overview of the available frequencies and the frequency auctions that were held are given in the appendices.

5.2 System Description

The following sections provide an overview of the different components that make up the Alvarion BreezeMAX network. Much of the detailed system information obtained during the project is confidential, hence specifications provided here are entirely based on information available to the public.

Figure 5.1 presents a high-level overview of the network components. The user is connected to the network via a subscriber unit (SU), also known as a mobile station. This may be a portable computer, a hand-held device, or a

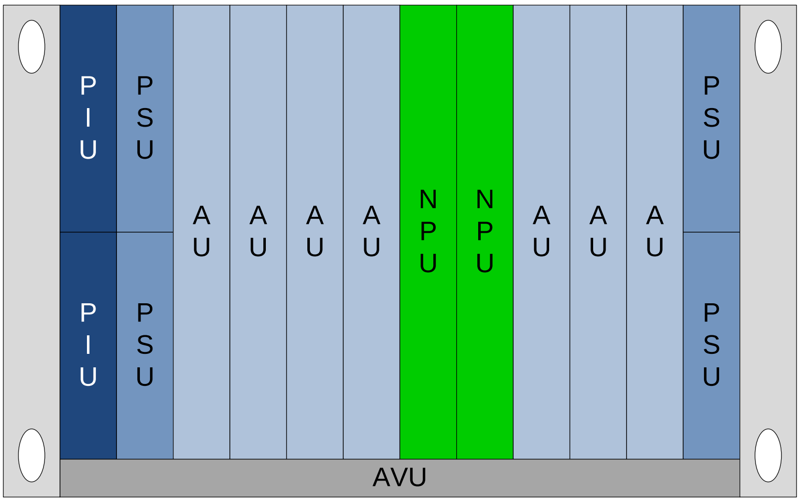

Figure 5.2: Base station architecture

residential gateway. The MS is connected to an antenna, which may appear in a number of configurations. Fixed WiMAX solutions commonly make use of a flat high-gain panel antenna mounted outside the premises. Indoor subscriber units with an integrated antenna may also be used. In the case of mobile stations, the antenna is embedded in the device. Both the antenna and the subscriber unit are referred to as customer premises equipment.

On the base station side, a number of antennas are present along with the base station components. The base station provides Internet access through the backbone network.

5.2.1 Base Station

The Alvarion BreezeMAX 2500 shelf edition was chosen as base station. It is certified against one of the fixed WiMAX certification profiles, but implements a number of functions that are included in the IEEE 802.16e-2005 amendment. In other words, it is said to be pre-mobile. The shelf version is modular and designed for high availability as all of its components support full redundancy.

The architecture of the shelf base station is illustrated in figure 5.2, and explained next. At the bottom, the 2U high air ventilation unit (AVU) includes a number of fans that provide airflow and cooling of the chassis. The 3U high power supply unit (PSU) and power interface unit (PIU) generate low-voltage DC output and feeds the entire base station with the necessary power. Up to four PSUs may be installed, providing a 3+1 redundancy configuration scheme.

The indoor access unit (AU-IDU) is 6U high and implements the IEEE 802.16 MAC and PHY layer functionality. Each AU-IDU controls a single sector and supports up to 4$^{\text{th}}$ order diversity using separation in either space, by polarisation, or both. The base station may support up to six sectors with a 6+1 redundancy scheme. The AU-IDU provides user traffic authentication and encryption along with handover management and QoS service flow management.

The network processing unit (NPU) is also 6U high and is the main controller of the base station. It is connected to the backbone through a dedicated Gigabit/Fast Ethernet interface and is responsible for aggregating traffic to/from all the AU-IDU modules. The NPU is also connected to a GPS for synchronisation[2]. In addition it provides in-band or out-of-band management functionality of the base station. The NPU performs VLAN encapsulation, traffic classification and tagging, and may implement ASN-GW functionality. There is support for two NPUs in a future 1+1 redundancy scheme.

In addition to the base station components already mentioned, each AU-IDU connects to one or more outdoor access units (AU-ODU) through distinct channels. The AU-ODU is a high power radio transceiver that amplifies and frequency-converts the signal from the AU-IDUs to the antennas and vice versa. Each AU-ODU connects to a single antenna element. For instance, if 2$^{\text{nd}}$ order diversity is wanted for a single sector, the following components are needed: A single AU-IDU, two AU-ODUs and two antenna elements. The elements can either be two separate single polarised antennas separated in space, or a single cross-polarised antenna. Similarly for 4$^{\text{th}}$ order diversity where 4 AU-ODUs and two cross-polarised antennas are needed.

The complete chain of components are outlined in figure 5.3, and table 5.1 lists some of the technical specifications. It is worth noticing that OFDM is used in the downlink, which means that the full bandwidth is available for each MS in turn. In the uplink, sub-channelisation is supported using OFDMA-16. This means that the system allows transmissions from several users simultaneously, with only a fraction of the sub-channels assigned to each user. This serves several purposes. First, any MS with poor link conditions can choose to transmit on a limited number of sub-channels instead of util-

[2]In a TDD system, it is imperative that all stations be synchronised in time and share the starting point of the frame. This is commonly achieved using a GPS clock at each base station. Network synchronisation via the backbone is also possible, but this functionality is not yet supported by Alvarion.

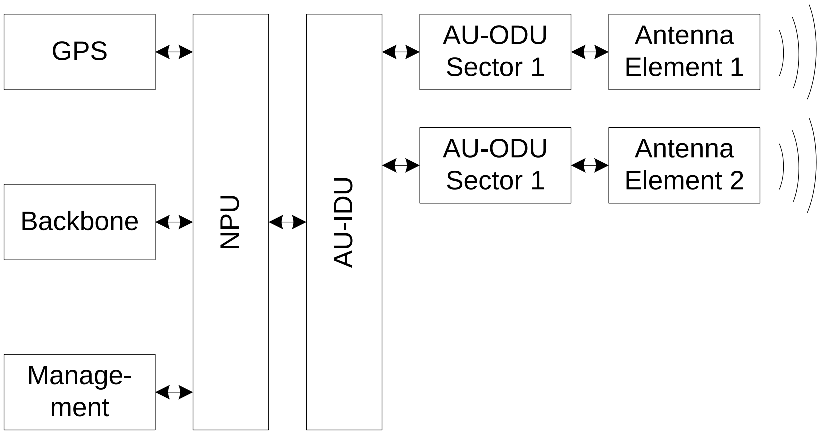

Figure 5.3: Base station configuration with 2^{nd} order diversity in a single sector. Note that two AU-ODUs and two antenna elements are used

ising the full bandwidth. By concentrating its power on fewer sub-channels, the maximum transmit power of the MS is increased, and higher availability is achieved. Second, by transmitting on fewer sub-channels, the mobile stations may conserve battery. And third, sub-channelisation better utilises the uplink capacity as multiple users may share the available bandwidth at the same time.

When using sub-channelisation in the uplink, while achieving higher availability, there is a great penalty in throughput. For instance, if a MS is transmitting from the cell edge using a single sub-channel, the throughput is reduced to 1/16 of what would have been achieved using the full channel. On the other hand, this leads to a 12 dB increase in the link budget.

Adaptive coding and modulation is also supported, where the modulation type and coding rate may be changed dynamically to accommodate varying channel conditions. When it comes to diversity schemes, STBC and CDD are supported in the downlink while MRRC is supported in the uplink.

5.2.2 Base Station Antennas

The system was set up with both no diversity and 2^{nd} order diversity in order to study any improvements in coverage and availability. Diversity was achieved by means of two different methods, namely separation in space and by polarisation. The former uses two vertical polarised antennas separated in space by at least 10 λ, while the latter uses a single dual-polarised antenna.

Parameter	Value
Frequency	2496-2690 MHz
Operation mode	TDD
Channel bandwidth	3.5, 5 MHz
Central frequency resolution	0.125 MHz
Output power	36 dBm maximum
Access method	OFDMA
Modulation	DL: OFDM, UL: OFDMA-16
	BPSK, QPSK, 16-QAM, 64-QAM
FEC	Convolutional coding: 1/2, 2/3, 3/4
Adaptive modulation and coding	Yes
Diversity schemes	DL: CDD and STBC, UL: MRRC
Management	Telnet, SNMP, and monitor port

Table 5.1: General technical specifications

Parameter	Value	Value
Type	Single	Dual
Frequency range	2300 - 2700 MHz	2300 - 2700 MHz
Gain	16.5 dBi	16.0 dBi
3dB azimuth beamwidth	60 degrees	65 degrees
Elevation beamwidth	7 degrees	8 degrees
Polarisation	Vertical	Dual-slant $\pm$ 45 degrees
Antenna port	N-Type, 50 Ohm	N-Type, 50 Ohm

Table 5.2: Antenna specifications

Both configurations were tested to investigate which would yield the best results.

The antenna specifications are listed in table 5.2. The single slant antenna is identical to the one that was used for signal propagation modelling.

5.2.3 Mobile Station

The CPE is comprised of three elements. The indoor unit (SU-IDU) comes in a number of configurations, however for this project a basic type was used. The SU-IDU is connected to the outdoor unit (SU-ODU) via a standard category 5 Ethernet cable. The SU-ODU contains all the active components such as modem, radio, and data processing capabilities. It also contains an integrated flat panel-antenna that is meant for fixed outdoor installation. The

Parameter	Value
Frequency range	2300 - 2700 MHz
Gain	8.5 dBi
Maximum input power	100 W
Elevation beam width	15 degrees
Polarisation	Vertical
Antenna port	N-Type, 50 Ohm
Weight	0.22 kg
Length	508 mm

Table 5.3: Antenna specifications of the mobile station

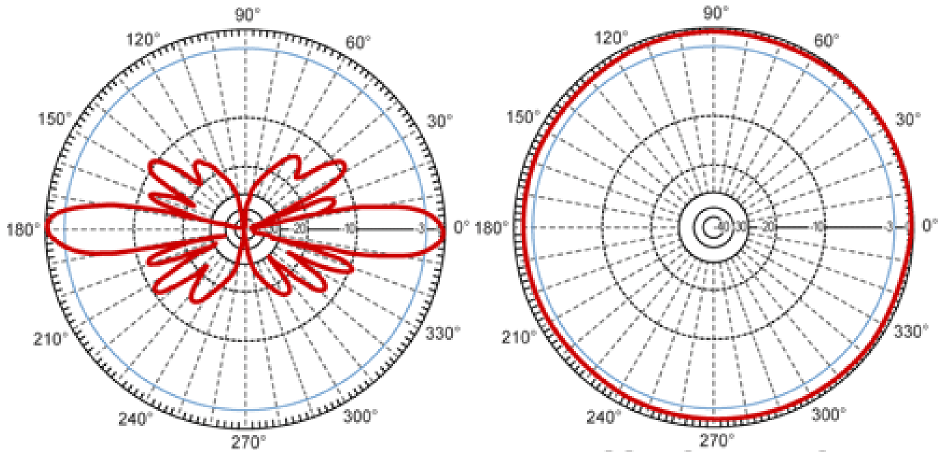

Figure 5.4: Mobile station antenna pattern

radio features Automatic Transmission Power Control (ATPC), where the power of transmission is dynamically adjusted depending on its distance to the base station. Its maximum output power is 19 dBm, which is considerably less than the base station.

In this project, an external omni-directional antenna was used instead of the flat panel-antenna that is integrated into the SU-ODU. The reason for this was that the CPE was to be mounted on a moving vehicle, and the need for an omni-directional antenna seemed obvious. The specifications for the antenna are listed in table 5.3, and figure 5.4 illustrates the radiation pattern.

5.2.4 Laptop Computers

Two laptops were used while testing the WiMAX network. The first one was a Dell Inspiron 500m with a 1.30 GHz Intel®Pentium®M processor and 1.25 GB of RAM. It was connected to the CPE at the client side, and most tests were run from this computer. The second one was a Dell Latitude D520 with a 1.66 GHz Intel®Core 2 Duo M processor and 512 MB of RAM. It was connected directly to the data port of the base station and was used both as endpoint during testing and for configuration of the base station. Both laptops were running Microsoft Windows XP Professional with Service Pack 2.

5.3 System Installation

In order to enable testing, the WiMAX base station equipment had to be installed at each of the two sites in turn. First, the system was set up at location BS1 and testing was carried out for the first half of the tram line. Next, the system was moved to location BS2, and testing for the remaining parts of the tracks were performed. At each site, testing was carried out both using no diversity and second order diversity in order to compare the two configuration schemes.

Only a short summary of the installation process is given here. More detailed information along with pictures describing the setup are given in the appendices.

5.3.1 Site BS1

The base station equipment was installed in a cabinet located at the rooftop of the Gunnerus library. A picture detailing the configuration is shown in figure 5.5. Fibre backhaul and 220VAC power was readily available on the roof. However, as the base station required -48VDC input power, a rectifier was also installed at the site to bridge between the PIU of the base station and the mains. In addition, a laptop computer was connected directly to the data port of the base station. This computer was used as endpoint for the throughput tests and for configuring the base station.

The outdoor GPS receiver was attached to a nearby pole while the antennas and AU-ODUs were mounted on poles and secured to the chimney. A total

Figure 5.5: Base station installed in the cabinet at site BS1

of three antennas were installed. Two vertical polarised antennas separated in space, and a single dual-slant antenna. Using this configuration, testing of no diversity and second order diversity is possible using both separation in space and by polarisation.

Figure 5.6 illustrates the antenna setup, where the left- and rightmost antennas are vertical polarised while the one in the middle is dual-slant. All antennas were mounted with an azimuth of 240 degrees, thus they were all pointing south-westwards. This is in accordance with the recommendations from the propagation modelling phase. No down-tilt was configured at the time.

5.3.2 Site BS2

The installation at this site very much resembled the configuration at the previous site. However, no fibre backhaul nor power were available at the rooftop. Because of this, the base station was not connected to the backbone while performing the tests. Instead, the same laptop as used previously was connected directly to the data port, from where both configuration and testing were performed. Power was provided via an extension cord from one of the offices below.

Figure 5.6: Antenna configuration at site BS1

The outdoor GPS receiver and antennas were mounted on poles and secured to the roof, as illustrated in figure 5.7. The initial antenna placement led to some problems during the preliminary testing, hence they were later moved closer to the edge of the roof. Only the two vertical polarised antennas were installed at this site, hence second order diversity using separation by polarisation was not tested from here. Both antennas were set up in accordance with the recommendations from the propagation modelling phase, with an azimuth of 300 degrees and no down-tilt.

5.3.3 Tram

The tram line company Gråkallbanen has five trams available for operation. Four of those are used in the daily service on a rotation basis. Unfortunately, only one of the trams had a 220VAC power socket installed. As the CPE required this voltage, only this tram could be used for testing. It is worth mentioning that the power received from this socket had a frequency of 100 Hz, as opposed to the normal 50 or 60 Hz found elsewhere. Luckily, both the laptop and the CPE worked well using this power source.

As illustrated in figure 5.8, the SU-ODU and the omni-directional antenna were strapped to the back of the tram. It was made sure that the antenna

Figure 5.7: Antenna configuration at site BS2

Figure 5.8: Antenna placement on the tram

was elevated above the roof of the tram to prevent unfortunate reflections leading to a possible degradation of the signal.

5.4 System Configuration

This section briefly describes the steps that were taken to configure and prepare the equipment for testing.

5.4.1 Base Station

The base station was initially configured by means of connecting to the NPU using Telnet. Each component of the BS was then configured in turn, starting with general parameters and then working through the other parts. In the following, the essential parameters that were defined are briefly explained. Some more detail is provided with respect to the services that were configured.

General Parameters

The switching mode was set to Ethernet-CS and ATPC was set to enabled. An optimal uplink RSSI value of -73 dBm was also configured, to prevent saturation of the base station and to provide optimal performance. This causes connected clients to dynamically adjust their transmission power, targeting an UL RSSI equal to the defined value.

Operator ID and cell ID are values that uniquely identify the operator and the given cell. For now, these were just left at their default values of 186.190.0 and 0.250 respectively.

Duplex mode was set to TDD with an UL/DL ratio of 50/50. The base station allows for a number of different ratios to be configured, depending on the expected load in the network. An important note to make is that this value is not changed dynamically to meet the needs of the users, even though such functionality is specified in the IEEE 802.16 standard. The reason for this is that when multiple base stations operate in TDD mode in the same area, they all need to be synchronised to prevent interference issues. Currently, there is no functionality implemented that allows the base stations to co-operate on such features.

The last set of general parameters that was configured were related to GPS synchronisation. Here, the one pulse per second (1PPS) signal from the external GPS receiver was enabled.

NPU

The management port and the data port were configured with appropriate IP addresses and subnet masks in order to enable connectivity with the backbone network. As recommended by the installation guide, in-band management was enabled. This means that remote management of the base station is made possible by transmitting data over a special VLAN to the data port. The dedicated management port on the NPU was used for configuration on site.

Radio Cluster

A radio cluster is a logical entity that represents one or more AU-ODUs which serve the same geographical sector. For the purpose of testing, one radio cluster was defined.

ODU

The outdoor access units were configured to use the 2590-2690 MHz frequency band, and their output power was set equal to 30 dBm, the minimum value. When diversity was used, both units were automatically set up to transmit at the same power level. Note that the value defined here is actually much less than what was used during propagation modelling, where an output power of 34 dBm was used.

AU

Only one indoor access unit was configured, as only one sector was being used for testing. The sector ID was set equal to its default value of 206. The maximum cell radius was defined as 20 km, and the operating bandwidth set to 5 MHz. Initially, no diversity was configured and only a single vertical polarised antenna was used in the sector. The frequency of operation was set to 2685 MHz.

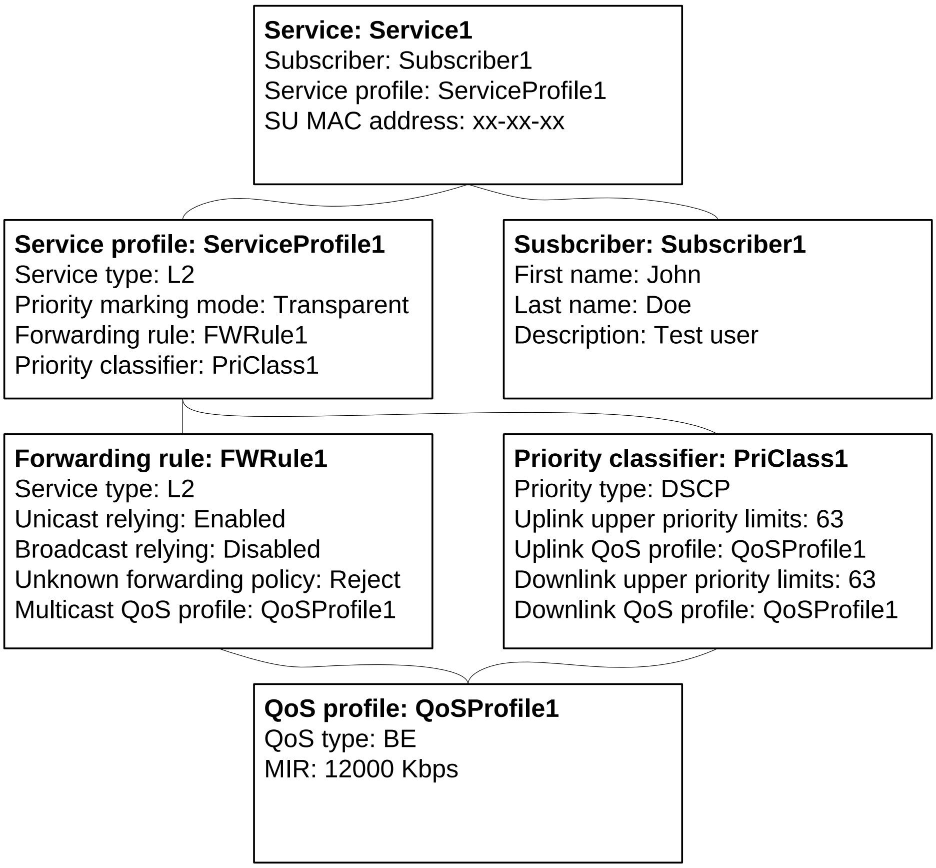

Figure 5.9: An overview of the defined service profile

SU

In order to keep things simple while testing, no RADIUS server was set up for authenticating the users. Instead, the MAC address of a single subscriber unit was added to the database of known units, allowing it to seamlessly connect to the base station.

Services

A service is a virtual connection between the subscriber's application and a network resource. It associates a given service profile with a subscriber's device, thus defining a set of forwarding rules and QoS parameters for each traffic flow. The service is implemented as an IEEE 802.16 connections, which was defined in section 2.6.1.

For the purpose of testing, a service named *Service1* was set up as described in figure 5.9. This service associated the service profile *ServiceProfile1* with subscriber *Subscriber1*. The service profile was set up with the L2 service type, a service which transport layer 2 Ethernet frames between the subscriber and the network resource to which it is connected. Transparent priority marking mode was defined, which causes the system to forwards frames without any changes to their DSCP bits. The service profile associated the forwarding rule *FWRule1* and priority classifier *PriClass1*.

The forwarding rule was defined for layer 2 traffic, where relying of unicast traffic was enabled and relying of broadcast and unknown traffic disabled. The *QoSProfile1* was associated with the relying of multicast traffic.

The priority classifier *PriClass1* was set up with DSCP priority type. Both uplink and downlink priority limits were set to their maximum values, meaning that all packets with such bits set would receive services as defined by the *QoSProfile1* profile.

Last, the *QoSProfile1* defined that its traffic flows should receive the best effort service type with a maximum information rate (MIR) of 12 Mbps.

5.4.2 Mobile Station

It is not necessary to perform many configuration changes to the mobile station, or subscriber unit. As long as its MAC address has been added to the list of permanent units in the base station, it nearly works out of the box. The only changes that were performed were adjusting its bandwidth of operation to 5 MHz and defining the start and end frequency ranges of scanning. The latter was done to reduce the time spent on the network entry process.

5.4.3 Laptops

The only changes that were performed on the laptops were related to IP configuration. Both laptops were set up with addresses on the same subnet in order for them to communicate across the wireless link. In addition, the laptop used on the client side was set up with an extra IP address which was used to connect to the SU-ODU for monitoring purposes.

Chapter 6

Network Testing

The outcome of the propagation modelling phase was a purely theoretical estimation of the achieved coverage. It was therefore necessary to verify the models against physical surveys, where a number of performance criteria were measured. This chapter starts with defining the different parameters that have been measured. Next, the different tests that were carried out to investigate the true performance of the WiMAX network are detailed. Following this, the actual results from the testing phase are presented.

6.1 Performance Criteria

The performance criteria that were measured are described in the following sub-sections. All physical parameters for both uplink and downlink were read directly from the CPE, while throughput parameters were gathered from the connected endpoints from where the tests were run.

6.1.1 Received Signal Strength Indication

The received signal strength indication (RSSI) is a measurement of the power level being received by the antenna. The value is commonly represented as a value in dBm, which is interpreted as the power ratio in decibel referenced to one milliwatt.

Section 8.3.9.2 in the IEEE 802.16 standard [7, 8] specifies that the MS shall obtain RSSI measurements from the OFDM downlink long preambles and report them via REP-RSP messages. Both mean and standard deviation statistics shall be reported in units of dBm, ranging from -40 dBm to -123 dBm.

6.1.2 Signal to Noise Ratio

The signal-to-noise ratio represents the power ratio of the desired signal to the background noise. Measuring the SNR yields a better estimation of the actual operating conditions of the system than the RSSI value, as noise is taken into account in addition to the received signal strength.

The IEEE 802.16 standard [7, 8] does not specify how the SNR ratio is to be measured, and leaves it to the implementers to decide. The MS shall report the value in units of dB via REP-RSP messages.

6.1.3 Modulation Rate

Based on the measured uplink and downlink RSSI and SNR values, the MS and BS decide upon which modulation rate to use for data transmission, and this may change dynamically from burst to burst. The different modulation rates supported by WiMAX were dealt with in section 2.5.5 of this thesis, and are not repeated here.

6.1.4 Number of Uplink Channels

Sub-channelisation in the uplink is supported by the WiMAX system, where a given MS may choose to focus its transmission power on a limited number of sub-channels in order to obtain higher availability from the edge of the cell. The technique was described in some detail in section 2.5.8 and section 5.2.1 of this thesis.

Measuring the number of sub-channels in use by the MS at any given time gives a clear indication of its current transmission capabilities. The value is measured as an integer from the set $\{1, 2, 4, 8, 16\}$, indicating the number of sub-channels in use.

6.1.5 Throughput

While this project focuses on measuring the physical attributes described above, throughput measurements are also important. The UDP protocol is the most appropriate protocol to use for determining the actual throughput of a system, as it does not require acknowledgements or re-transmissions of lost packets. However, TCP is the protocol used by the majority of Internet applications, hence measuring the throughput of this protocol yields results that are more realistic when it comes to web-browsing and file downloading.

There is one great drawback of using TCP over a wireless link. If the protocol experiences packet loss, which happens all the time when it comes to wireless transmissions, it wrongly assumes the network is congested and reduces its transmission rate. TCP then rebuilds its transmission window relatively slowly, leading to a sawtooth throughput curve. In other words, when measuring TCP throughput over a wireless link, the actual results may be too pessimistic.

Testing was carried out using IPERF [42], and because of time constraints, only TCP throughput in the downlink has been measured. It is important to realise that TCP is in fact a two-way protocol that requires acknowledgements to be sent in the opposite direction of the data stream. Because of this, TCP throughput relies on the channel conditions in both directions.

6.2 Overview of Testing

The first set of tests were carried out with the base station equipment installed at site BS1, with no diversity configured. This means that only a single vertical polarised antenna was used at the base station side. Initially, a set of preliminary tests were performed from a few random locations within the coverage area. The purpose of these tests were to verify that the MS could successfully connect to the BS and that the signal strength and throughput measurements were according to what was expected. The tests also provided knowledge about the coverage area, to make sure it was in accordance with what was predicted by the propagation model. Important performance criteria in this phase were uplink and downlink RSSI and SNR, and TCP throughput.

After successfully completing the preliminary testing phase, live tests on the tram were carried out. Performance measurements such as uplink and downlink RSSI and SNR, modulation order, transmit power, and the number of sub-channels in use, were continuously gathered from the MS for the entire journey from start to end. From a connected GPS, altitude, speed, and position data were also be measured. Simultaneously, the throughput of a TCP connection was benchmarked. The data from all these tests were then linked in order for the results to be analysed later.

The tests on the tram were carried out both with and without second order diversity. The reason for this was to investigate the difference in diversity orders with respect to availability and signal quality. First, a couple of tests were run with no diversity configured. Next, second order diversity using two vertical polarised antennas separated in space was tested. And last, second order diversity using a single dual-slant antenna was investigated. Each test was carried out several times in order to mitigate intermittent errors.

Afterwards, the base station equipment was moved to site BS2 and the complete set of tests were carried out once more on the remaining length of the tram line.

Note that during all these tests, only a single CPE was used in the network. This means that the full capacity of the system was available to this unit at any one time.

6.3 BS1: Preliminary Tests

The first set of tests were performed from a few random locations along the tram line. The base station was configured with no diversity, hence a single vertical polarised antenna was used in the sector. A total of eight different locations were visited, and both signal quality measurements and throughput measurements were carried out. Additionally, a ninth test was undertaken by driving along the tram line to investigate the performance while in motion.

Figure 6.1 shows the locations of the eight fixed points from where measurements were taken. As can be seen, all points lie in the near vicinity of the track. The location of the base station is also marked. The numbered stars are to be interpreted as follows. A white star indicates a location where no connection to the base station was achieved, while a black star indicates where a connection was successfully set up. A black and white star indicates that a connection was successfully set up, but the data throughput was virtually zero. Important physical performance data are outlined in table 6.1.

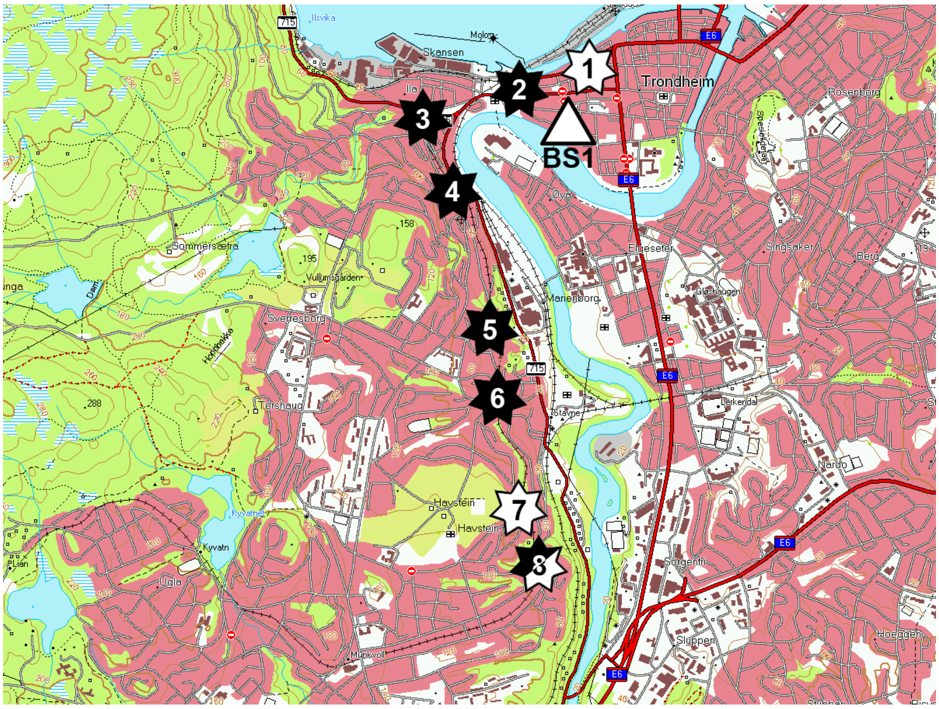

Figure 6.1: Overview of preliminary testing from BS1. White stars indicate no connection. Black stars indicate successful connection. The black and white star indicates successful connection to the BS, but with a signal too weak for useful communication.

Location	RSSI [dBm]		SNR [dB]		TX Power [dBm]	Distance [km]
	DL	UL	DL	UL		
1						0.30
2	-88	-85	12	16	18	0.39
3	-73	-73	28	27	19	0.95
4	-58	-73	33	26	-2	0.89
5	-90	-89	13	13	13	1.40
6	-81	-77	22	25	17	1.90
7						2.50
8	-100	-96	3	5	17	2.90

Table 6.1: Physical performance results from preliminary testing from BS1

6.3.1 Comparison with Propagation Modelling

The first thing to investigate is whether or not the measured coverage coincides with the results from the propagation modelling phase. Comparing figure 4.7 with the results from table 6.1 yields the following.

Location 1 lies outside the expected coverage area, hence the lack of signal here is in accordance with the modelling. At location 2, the measured signal strength was a bit weaker than what was expected. The reason for this is probably the amount of large buildings blocking the propagation between the base station and the client. Because no building height file was used during modelling, the predicted coverage was a bit optimistic.

At location 3 and 4, the measured signals were in accordance with the modelling. However, at location 5 there was severe degradation. The modelling results predicted a very strong signal here, -76 dBm or better, as represented by the green/yellow area between point 7 and point 8 in figure 4.7. Instead, the measured DL RSSI was only -90 dBm. The reason for this is not clear, but the low strength may be related to both dense vegetation and a few houses blocking a line of sight propagation to the base station. At location 6 the signal was as expected.

According to the propagation modelling, location 7 is located right outside the coverage area. The reason for not obtaining connectivity from here seems clear, as there was dense vegetation blocking the propagation from the base station. At location 8, the measured signal was also more or less in accordance with the modelling. Here, the signal was too weak for any data to go through, measuring -100 dBm in the downlink.

Even though the number of visited locations is quite low, the conclusion so far is that the modelled results to a large degree coincides with the measured ones. The reasons for achieving less signal strength than expected at some of the locations may be caused by the following:

- Testing was conducted from random locations without performing micro positioning of the antenna[1]. According to [43], micro positioning leads to severe increase or decrease in signal strength. Moving the antenna a few meters could easily impact the received signal strength by 5 to 10 dB.

[1] The signal strength varies considerably within a small area due to multipath fading, as discussed in section 2.7.2. In order to find an optimal spot where the signal is stronger, the antenna may be moved around a bit. This is known as micro positioning.

- No building height file was used during modelling, hence the predicted coverage in the city may be too optimistic.
- The clutter layer used during modelling had a resolution of 100 metres, which is relatively coarse. Because of this, all the small variations in ground occupancy were not taken into account by the model.
- The BSs were configured with a transmit power of 34 dBm during modelling. In contrast, the output power was only 30 dBm while performing field testing, a difference of 4 dBm.
- The base station antenna gain was set equal to 17 dBi during modelling, while it was actually 16.5 dBi during testing, a difference of 0.5 dBi.
- The CPE antenna gain was set equal to 6 dBi during modelling, while the real antenna had 8.5 dBi of gain.
- Cable and connector losses at both the base station and the CPE was set equal to only 1 dB during modelling. In practice, these numbers were probably a bit higher.

6.3.2 RSSI as a Function of Distance

The received signal strength decreases with distance, as the losses defined by the radio propagation models discussed in section 4.3.1 increases. In order to calculate the expected received power, a simple link budget must be derived:

$$\begin{aligned} ReceivedPower_{dBm} &= TransmittedPower_{dB} \\ &+ TransmitterAntennaGain_{dBi} \\ &+ ReceiverAntennaGain_{dBi} \\ &- CableandConnectorLoss_{dB} - PathLoss_{dB} \\ &= 30dB + 16.5dBi + 8.5dBi - 4dB - PathLoss_{dB} \end{aligned} \tag{6.1}$$

In this budget the transmitter power is set equal to 30 dB and the antenna gain at the base station side is equal to 16.5 dBi when no diversity is being used. The receiver antenna has a gain of 8.5 dBi and the sum of cable and connector losses at both sides is equal to 4 dB. The parameters used in the path loss models are the following: Frequency f equal to 2680 MHz, H_b equal to 20 metre, and H_m equal to 2 metre. In other words, the parameters used in the link budget are as close as possible to reality.

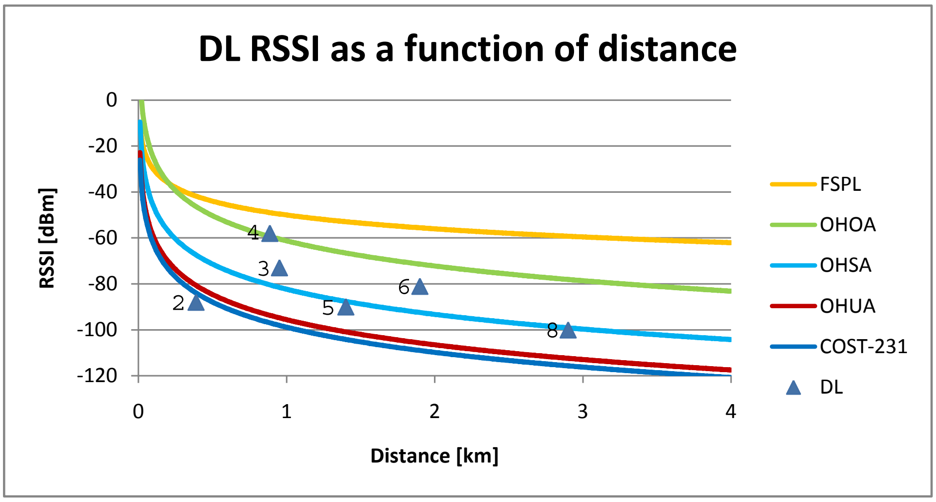

Figure 6.2: DL RSSI as a function of distance from BS1 compared to path loss models

Figure 6.2 is a graphical representation of the DL RSSI values plotted as a function of distance from the base station. The five path loss models discussed earlier are also shown for comparison. The left-most measurement, labelled 2, was taken well inside the city, and coincides with the model for urban areas. Looking closely at the graph, is can be seen that the signal strength at this location is actually a bit less than what would be expected. The cause of this may be that operating in the side-lobe of the antenna leads to some degradation. The next measurement, labelled 4, was taken with a clear line of sight to the base station, hence it falls on the curve for open areas. The rest of the measurements more or less fall in between the models for open and sub-urban areas. This makes sense as the area between the base station and the different test locations is indeed a mix of such environments, and thus favourable with respect to radio propagation.

The values for uplink RSSI are more difficult to plot, as the MS uses ATPC to adjust its output power. This means that the MS targets the optimal UL RSSI value as defined in section 5.4.1, thus its transmission power is not constant. The result is a much smaller dynamic range in the values for uplink RSSI compared to downlink RSSI, as indicated by the data in table 6.1.

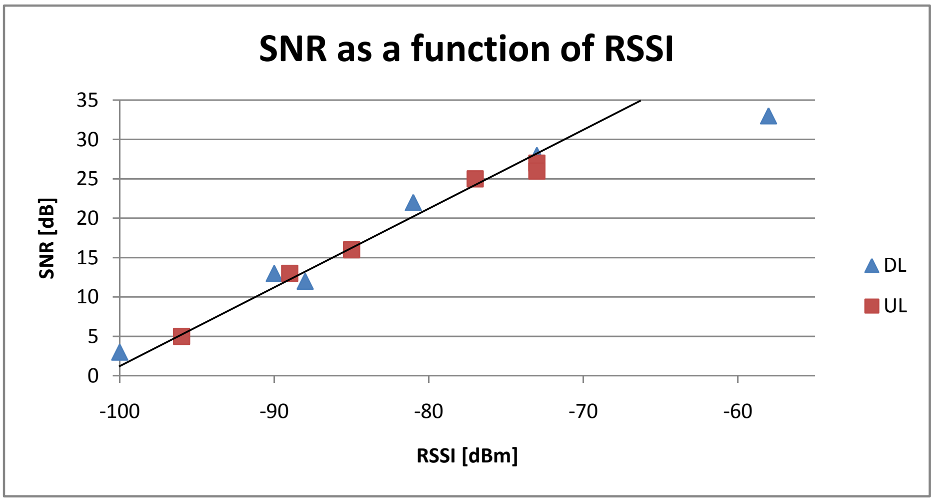

Figure 6.3: SNR as a function of RSSI for both uplink and downlink

6.3.3 SNR as a Function of RSSI

The SNR is a measure of the ratio between the received signal and the noise in the channel. If the noise is constant and interference is absent, the relationship between SNR and RSSI should result in a linear graph. The resulting plot is presented in figure 6.3, and even though the number of points is relatively low, it can be seen that the two are highly correlated. An increase of 1 dB in SNR results in an increase of 1 dB in RSSI.

The only exception is the right-most point, which falls outside the linear regression. According to [43], this is caused by the MS, which cannot measure such high values for SNR.

6.3.4 TCP Throughput as a Function of RSSI

The throughput of a TCP downlink connection was measured at all locations, while the uplink was only measured at locations 4 to 8. The averaged numerical results are presented in table 6.2, and graphically illustrated in figure 6.4.

As can be seen, the throughput increases with increasing RSSI. This relationship makes sense, as higher RSSI allows for higher orders of modulation, hence more bits may be transmitted per symbol time. The highest throughput measured was 6408 Kbps in the downlink at -58 dBm. However, this bitrate is only $6408/9072 \approx 70\%$ of what is theoretically possible. Protocol

Location	TCP DL throughput [Kbps]	TCP UL throughput [Kbps]
1	N/A	N/A
2	155	N/A
3	2016	N/A
4	6408	5973
5	380	1352
6	1677	3055
7	N/A	N/A
8	0	0

Table 6.2: Average TCP throughput from preliminary testing from BS1

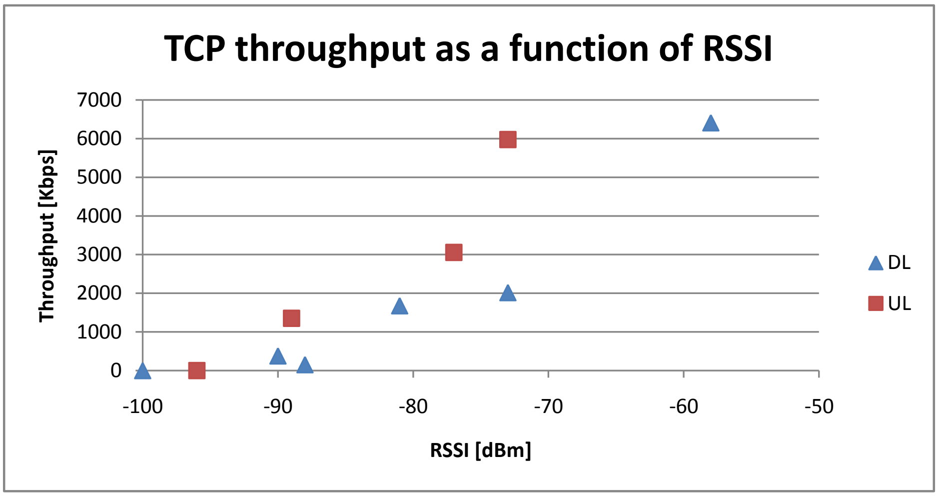

Figure 6.4: TCP throughput as a function of RSSI

overhead and capacity reserved for management traffic may account for some of this difference.

The downlink result from location 2 is relatively low, achieving only 155 Kbps. This was due to the TCP send and receive windows, which caused underutilisation of the wireless link. When using TCP to measure throughput, it is important to understand the inner workings of the protocol, hence it is explained in some detail in the following.

In order to limit the amount of data that can be sent by the transmitter and provide the receiver side with flow control, TCP uses a sliding window. For a given direction of the full-duplex TCP connection, the sender maintains a send window and the receiver maintains a receive window. These windows control the amount of data that can be in transit simultaneously, and is a mechanism that prevents receiver saturation. The size of the window field in the TCP header is 16 bits, allowing either side to advertise a maximum receive window of 65,535 byes without window scaling. Using window scaling, the window may be set even larger.

In order to optimise TCP throughput, the logical pipe between the transmitter and receiver should be full of data at any one time. The capacity of the logical pipe is known as the bandwidth-delay product (BDP), which can be calculated using the following formula:

$$BDP = bandwidth * RTT \tag{6.2}$$

where bandwidth is measured in bits per second and the round-trip time (RTT) in seconds. Pipes that are fat (high bandwidth), and long (high RTT) have the highest BDP. Examples of high BDP links are transmission paths across satellites or via long distance fibre links, both which have high bandwidth and long delays. In this WiMAX pilot, the maximum bandwidth using a UL/DL ratio of 50/50 and highest modulation is typically 9 Mbps in each direction, according to product specification. The RTT has been measured using the ping tool and averaged to 50 ms. In this case, the BDP is equal to $9Mbps*0.050s = 450Kb = 56KB$. If the receive window is less than 56 KB, the link will be underutilised. Because of this, it is important that the receive windows at both sides of the communication link are increased to a value larger than the BDP. However, if the window is set too large, it may fill up the available memory at the receiver.

Thus, for the first measurements, the TCP windows were set too low, hence the measured throughput was much less than what could be expected. For location 3 and onwards, the windows were set to a much larger value in order to mitigate this problem.

Referring once more to table 6.2, it can be seen that the throughput in the uplink is actually much higher than in the downlink. The difference amounts to more than three times as high at location 5, and almost twice as high at location 6. Consulting the uplink and downlink RSSI and SNR values from table 6.1 yields that there is not much difference at location 5. At location 6 there is a difference of 3 dB in SNR, which could explain why the uplink achieves better data rates than the downlink.

6.3.5 Driving Along the Tram Line

In addition to the eight fixed points visited during the preliminary testing phase, a ninth test was conducted by driving along the tram line while performing live measurements. The idea was to investigate how the system would respond to a moving vehicle, trying to resemble the tram as much as possible. An overview of the track driven is given in figure 6.5. As can be seen, the track runs parallel to the tram line at all times. Average speed during testing was 20 kmph and the length of the run was 1.8 km.

Physical performance data was measured every second for the entire journey. In addition, other parameters such as speed, altitude, and geographic coordinates were gathered from a connected GPS. The result yielded more than 300 distinct measurements pin-pointed in both time and location. Note that no throughput measurements were taken during the run.

Figure 6.6 shows the DL RSSI values plotted as a function of the distance between the MS and the base station. The five different path loss model are also included for comparison, with parameters identical to what was used before. It is interesting to note that most measurements fall in between the Okumura-Hata models for open and sub-urban areas. This makes sense, as the area between the MS and the BS is indeed very open, with some dense vegetation and buildings blocking a direct line of sight. Looking at figure 4.2 confirms this, where the track driven coincides to a large degree with the solid black line to the left in the picture. It is also interesting to note that some measurements approach the FSPL, which indicates a clear line of sight to the base station.

As more measurements were taken, a new plot of the SNR as a function of RSSI was made. The result is graphically illustrated in figure 6.7, and interpreted as follows. There is a clear tendency that an increase in SNR

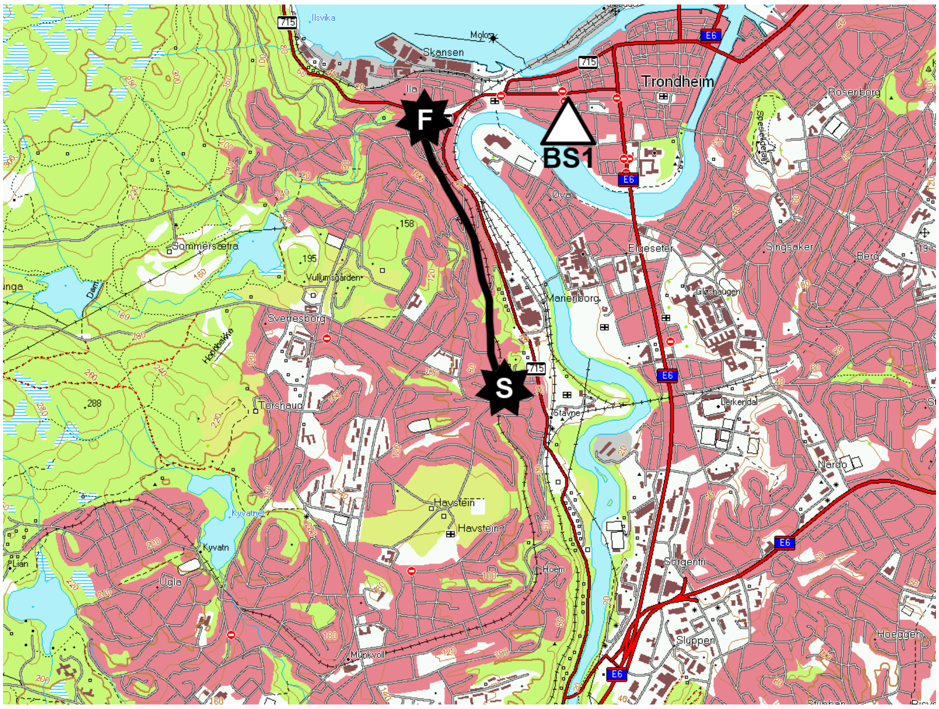

Figure 6.5: The track used for the ninth test during preliminary testing from BS1. The start and finish points are marked with black stars.

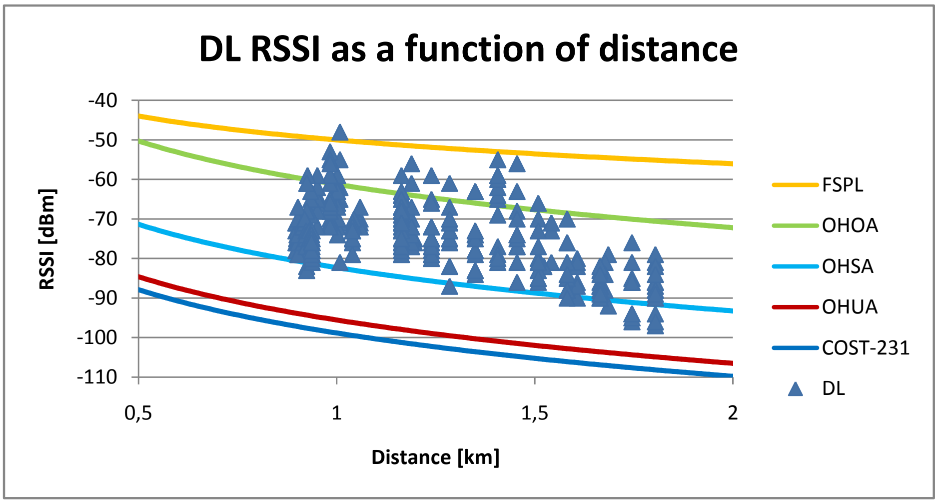

Figure 6.6: DL RSSI as a function of distance from BS1 compared to path loss models

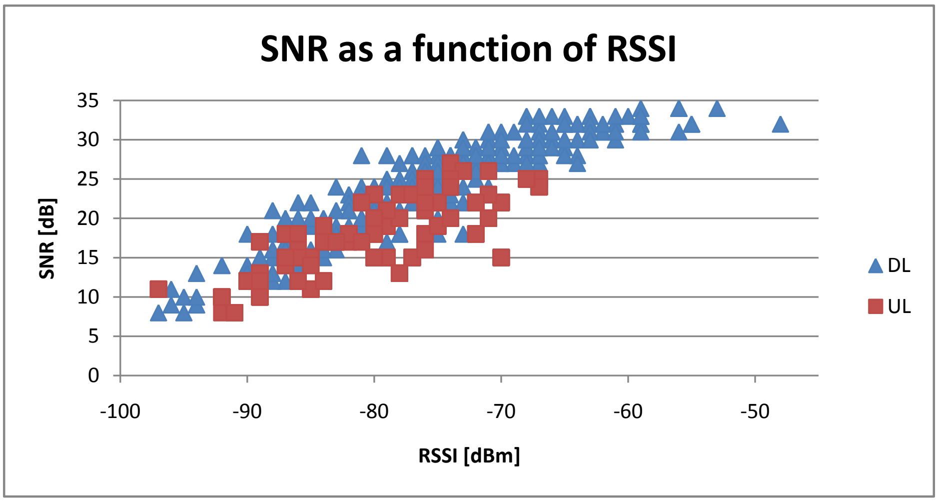

Figure 6.7: SNR as a function of RSSI for both uplink and downlink

yields a similar increase in RSSI. In other words, there is a linear relationship between the two values. However, the variance is much greater than what was found during the tests from the fixed locations. The reason for this is that the noise in the channel is not fixed, meaning that for identical values of RSSI, the SNR differs.

It is also interesting to note that the dynamic range of the UL RSSI values are much smaller than the DL RSSI values. This fact was noted earlier, as the MS uses ATPC to adjust its output power in order to avoid saturating the receiver. The base station was configured with an optimal uplink RSSI value of -73 dBm, hence the MS targets this number. Looking at the graph, it can clearly be seen that most UL RSSI values fall below this threshold.

Last, the right-most points seem to level off from the linear relationship. As commented on earlier, this illustrates very clearly that the MS cannot measure such high values for SNR.

6.4 BS1: Testing on the Tram

The next set of tests were carried out on the tram, where several journeys were made from start to end while performing measurements. Some of the more interesting observations are presented in the following sub-sections.

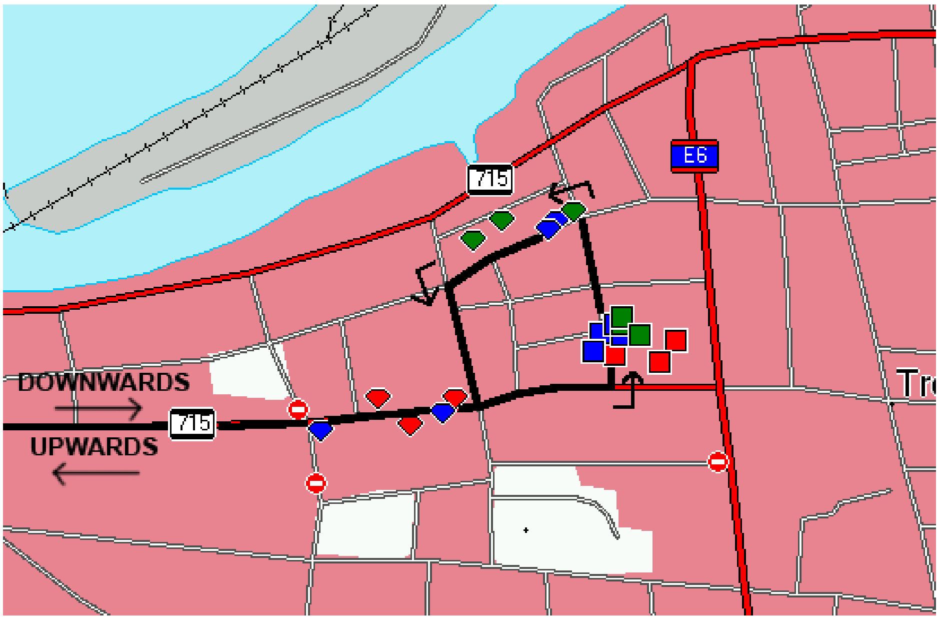

Figure 6.8: Signal reach at the end station in the city centre. Diamonds indicate where the signal was picked up on trips upwards, squares where it was lost on trips downwards. Red means no diversity, blue and green means 2nd order diversity using VV or X respectively.

6.4.1 Signal Reach

What impact does diversity have on signal reach? According to theory dealt with in section 2.7.2, 2nd order diversity using STBC in the downlink and MRRC in the uplink should indeed increase the coverage. In order to verify this, the points where the signal was picked up and lost were plotted on a map. Several trips were made both upwards and downwards, thus the number of measurements taken is relatively high.

The results from the end station in the city centre are shown in figure 6.8, where the tram line is marked with a black, solid line. The diamonds indicate where the signal was picked up on the trips upwards, and the squares indicate where the signal was lost on the trips downwards. The red colour means no diversity, blue and green means 2nd order diversity using two vertical polarised antennas (VV) or a single dual-slant antenna (X) respectively.

When it comes to picking up the signal, there is a clear difference between no diversity and 2nd order diversity, even though the difference is rather small. Using no diversity, the results coincides very well with what the propagation

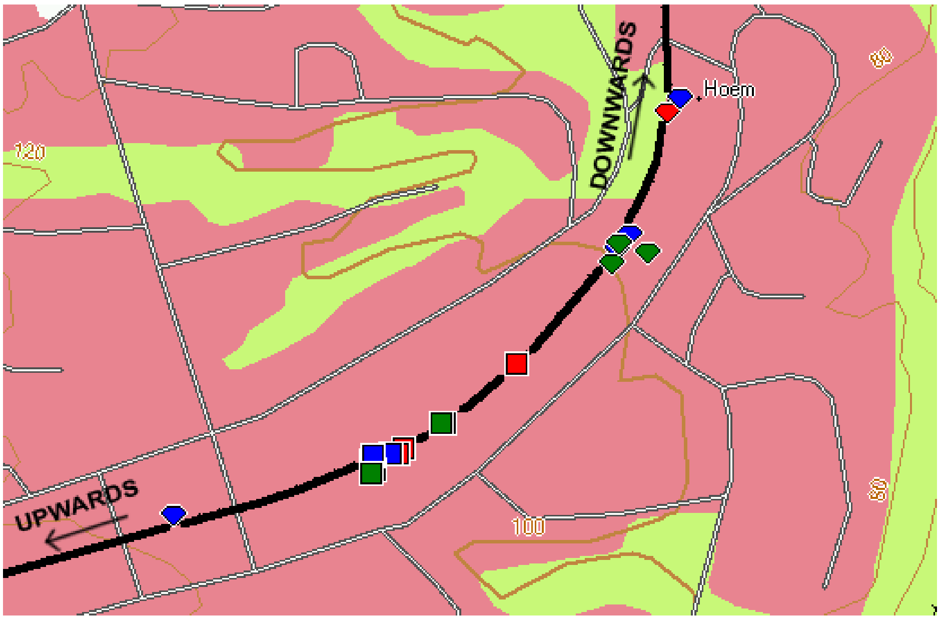

Figure 6.9: Signal reach at Hoem. Diamonds indicate where the signal was picked up on trips downwards, squares where it was lost on trips upwards. Red means no diversity, blue and green means 2^{nd} order diversity using VV or X respectively.

model predicted. Comparing the measurements to figure 4.7, it can be seen that the three red diamonds fall right on the line indicating the start of the coverage area. When using diversity, the signal is picked up earlier in all but two occasions, as indicated by the two blue diamonds among the red ones. It is interesting to note that the group of both the red and the green diamonds are densely spaced, indicating a high correlation between subsequent runs. The difference in meters between the two groups is a bit less than 200 metres.

Loosing the signal happens at the exact same spot with or without diversity in use. The inaccuracy of the GPS reception accounts for the small variance in exact location of the squares. In this case, diversity does not seem to have an effect at all.

From these measurements it is realised that the current base station configuration using a single antenna pointing south-westwards is not optimal with regards to covering this particular area. As indicated both from the propagation modelling phase and these measurements, it is necessary to implement another sector at site BS1 to provide adequate coverage at the start of the tram line.

Figure 6.9 shows the results from Hoem, which lies in the middle of the tram line, approximately 4.5 km from the starting point. Once again, the diamonds indicate points where the signal was picked up on the journeys downwards, while the squares indicate signal loss on journeys upwards. First, only a single red diamond is shown in the figure. The other two are located further north, closer to the base station, indicating even worse signal reception. It is interesting to note that the three green diamonds are all located at the very same spot. This indicates a high correlation between different runs, and clearly states the exact point from where the signal strength is sufficient for communication. The blue diamonds seem to be scattered over a larger area. One point coincides with the green group, while another is located a hundred meters further north. The left-most blue diamond seems to be misplaced. However, the CPE actually managed to connect to the base station from this point because the tram stopped for a long time here. When the tram started moving again, it soon lost the connection before re-connecting again close to the green group.

It is evident that there is a certain difference between using no diversity and 2nd order diversity when it comes to picking up the signal. Especially separation by polarisation shows the most promising results, while separation in space fluctuates more. When using no diversity, the tram has to move a few hundred meters further north before obtaining a steady connection.

When it comes to signal loss, little or no difference was measured. The loss of the signal happened at the exact same spot each time, which was in accordance with what was predicted in the propagation modelling phase.

6.4.2 RSSI as a Function of Distance

When 2nd order diversity is being used, the base station transmits from two antenna elements simultaneously, hence the EIRP is doubled. Because of this, the client should notice both a stronger signal and increased coverage. In the previous sub-section it was concluded that diversity led to a small improvement in coverage area. In this section, the DL RSSI values obtained from using no diversity and 2nd order diversity are investigated.

Comparing the DL RSSI values measured by the client for each individual run led to some initial problems. Even though measurements were gathered each second, linking the data based on the time since departure made no sense as the tram did not pass the same locations at identical points in time. Sometimes the tram would stop for red light, other times it would stay at a stop for a long time. Hence, because the journey from start to end varied in time[2], the data could not be linked based on this criteria.

Instead, the data could be linked by distance. As each measurement was accompanied by a set of coordinates, the distance to the base station from each point along the track could easily be calculated. This method was tried initially, but some problems arose because the tram would pass through several different points with an equal distance to the base station. As a result, the measured DL RSSI values could not be directly linked to this distance.

The track length was used as reference instead. Because the tram would travel the exact same distance each time, the measurements could be linked by the distance from the starting point. However, two small problems arose. First, as the location received from the GPS was inaccurate, the measured length travelled by the tram varied a bit from run to run[3]. And second, the GPS only plotted a point for every 37 meter on average, while the physical performance parameters were read every second. Because only one reading was linked to each GPS point, much of the information gathered between any two points were not taken into account when the data was analysed. However, obtaining data for every 37 meter on average along the track yield results that are fairly accurate, and the trend line is clearly seen.

The geographic coordinates received from the GPS were on the decimal degree form. This means that the latitudinal and longitudinal values were converted from degrees, minutes, and seconds to degrees and a decimal value for the minutes and seconds, thus making calculations easier. However, in order to calculate the distance in meter between any two points, the degrees and minutes had to be converted to the metric system. This was straightforward using a set of formulae included in the appendices. After finishing the calculations, the physical performance measurements could easily be linked to their distance along the track.

A graphical illustration of six DL RSSI measurements are presented in figure 6.10, where each measurement has been plotted as a function of distance from the starting point of the tram line. In this case, no diversity was used. As can be seen from the figure, there is a high correlation between each of the runs. Because of this, the data from each set of tests have been averaged to produce

[2]Based on 19 measurements, the moving time of the tram varied from 18:20 to 21:24, a difference of three minutes. Its average moving time was 19:58. The entire trip time, including all stops, varied from 23:56 to 26:33, a difference of 5 minutes. Its average trip time was 23:56. On average, the tram stopped for four minutes during each journey.

[3]Based on 16 measurements, the minimum distance measured was 8.72 km and the maximum was 8.91 km, a difference of 190 meters.

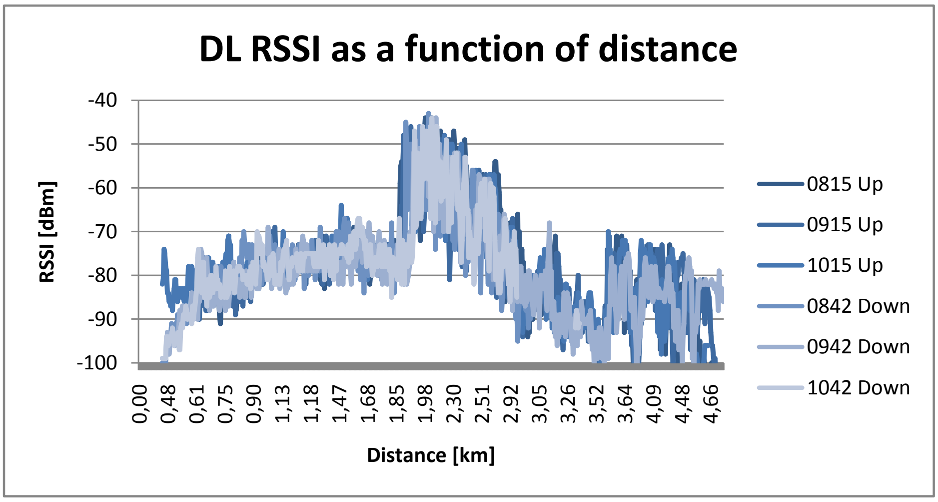

Figure 6.10: DL RSSI values for six consecutive runs without diversity from BS1

Diversity	DL RSSI [dBm]		UL RSSI [dBm]	
	Upwards	Downwards	Upwards	Downwards
None	-78	-77	-79	-80
VV	-74	-78	-79	-82
X	-77	-79	-81	-83

Table 6.3: Average physical performance measurements from BS1

clearer and more readable graphs. By doing so, sources of intermittent errors are also mitigated.

The average DL RSSI results from 18 journeys on the tram are presented in figure 6.11. The red line illustrates the results from using no diversity, while the blue and green lines represent diversity using separation in space and by polarisation respectively. Table 6.3 presents average DL and UL RSSI values for all three scenarios, separated in direction upwards and downwards.

The conclusion from these measurements is that diversity does not seem to have a great impact on the physical performance parameters. While 2nd order diversity using separation in space indeed achieves an amazing 4 dB increase in DL RSSI compared to using no diversity when travelling upwards, separation by polarisation only performs 1 dB better on average. The same parameter when travelling downwards actually performs worse with diversity, lowering the average by 1 and 2 dB respectively. The averaged UL RSSI values do not benefit from diversity at all, in both scenarios the measured

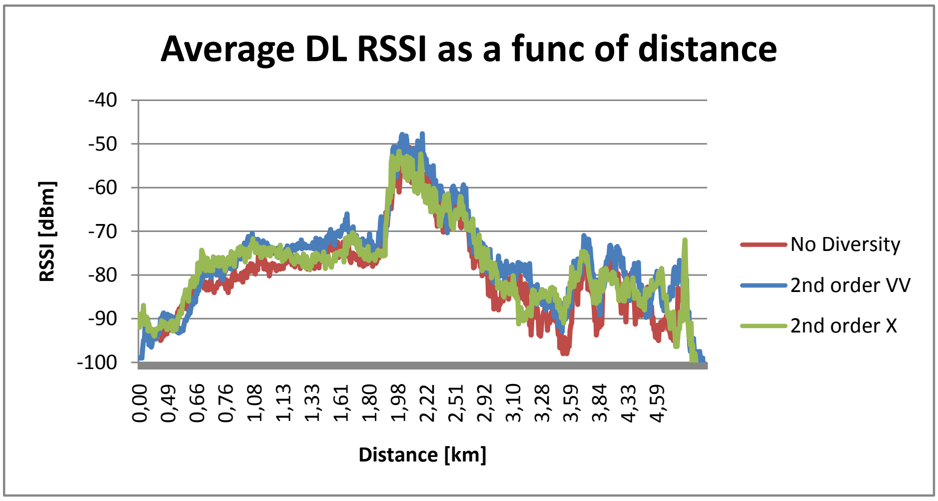

Figure 6.11: Average DL RSSI values from 18 consecutive runs with and without diversity from BS1

values are either identical or lower than when not using multiple antennas.

Generally, a better signal is achieved in both uplink and downlink when travelling upwards. The antenna placement on the tram may be accountable for some of this variance. As described in section 5.3.3, the antenna was strapped to the back of the tram and elevated above the roof. When travelling upwards, the antenna faces the base station most of the times. When going in the opposite direction, the roof may actually block or reflect parts of the signal energy, thus causing a degradation.

Last, figure 6.12 shows the DL RSSI values plotted as a function of the distance between the tram and the base station. The five different path loss model are also included for comparison. The plot is based on six journeys along the tram line, three in each direction. No diversity was used, and a total of almost 3300 measurements were taken. It is interesting to note that almost all points within half a kilometre of the base station fall at or below the Okumura-Hata model for urban areas. There are two obvious reasons for this. First, this part of the tram line runs through the inner city, where high buildings block signal propagation. And second, because this area is covered with the back-lobe of the base station antenna, the signal quality here is not very good. At one kilometre, there is a peak in performance, where all points perform better than the model for open areas. Here, there is a clear line of sight to the base station, which makes propagation favourable. All points further away fall in between the models for open areas and sub-urban areas.

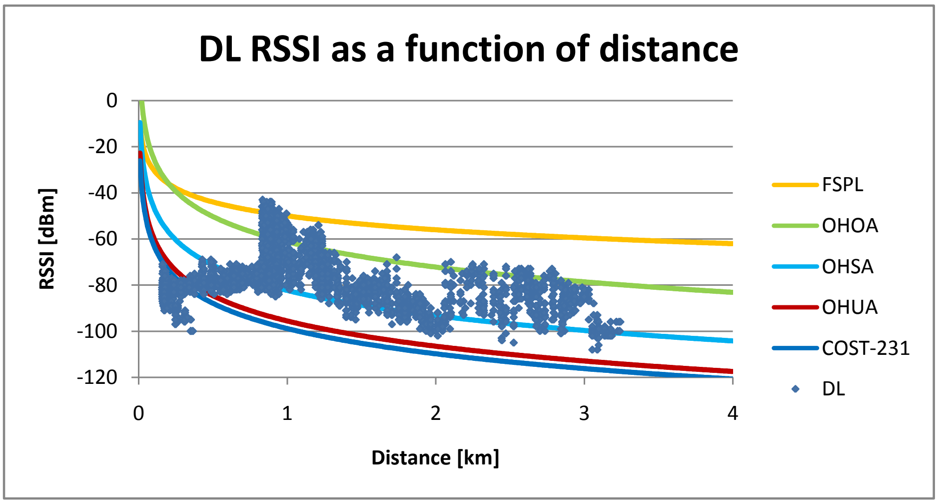

Figure 6.12: DL RSSI as a function of distance from BS1 compared to path loss models. No diversity was used

Diversity	TCP Throughput [kbps]		
	Upwards	Downwards	Average
None	200	287	243
VV	313	375	344
X	272	282	277

Table 6.4: Average TCP throughput measurements from BS1

This coincides well with was measured during the preliminary testing phase, described in section 6.3.5.

6.4.3 TCP Throughput

The throughput of a TCP DL connection was measured for each run in order to investigate if speed or diversity had an impact on the achievable data rate. Figure 6.13 shows the averaged results separated by diversity order, and table 6.4 contains numerical data separated by both diversity order and direction of travel. The results are based on three round-trip journeys using no diversity and six journeys using each of the diversity configurations, yielding a total of 18 runs.

As indicated by the graph, the throughput increases slowly from the start, reaching an average of 400 kbps for all three diversity configurations after about one kilometre along the tram line. This first part of the track was ear-

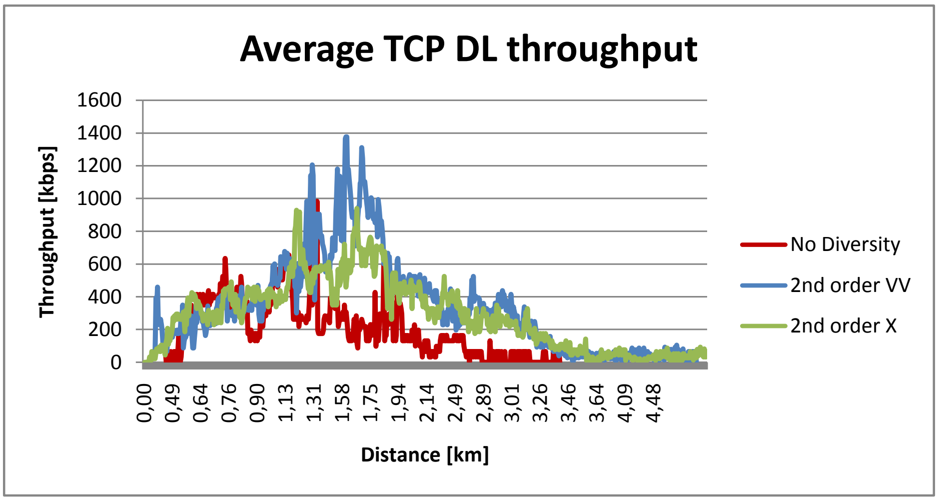

Figure 6.13: Average TCP DL throughput from BS1 as a function of distance along the tram line. The red line illustrates no diversity, the blue and green lines represent 2nd order diversity using separation in space and by polarisation respectively

lier identified as an area with varying signal quality, hence the low throughput can be explained. For the 2nd order diversity configurations, the throughput increases to an average between 400 and 1400 kbps between the one and two kilometre marks. Refer to figure 4.1 for a map of the tram line. From the two kilometre mark and onwards, the throughput decreases steadily, reaching a very low value at 3.5 kilometres. During the next kilometre, the throughput drops to virtually zero. When using no diversity, the a throughput of zero is reached at the 2.5 km mark.

Generally, the average throughput when using no diversity is much less than for the diversity configurations. This observation is supported by the numerical values from table 6.4, which states that 2nd order diversity using separation in space achieves $344/243 - 1 \approx 42\%$ better results than using no diversity, while separation by polarisation achieves only $277/243 - 1 \approx 14\%$ better. There is also a noticeable difference between upwards and downwards journeys, where the latter performs better. However, this is not the case for 2nd order diversity separated by polarisation, which performs just as well in both directions. This was not as expected, as both the UL and DL RSSI values were measured to be higher for upwards than downwards journeys. There is no clear explanation to why there is such a difference.

During testing, it was noted that the throughput was very dependent on the

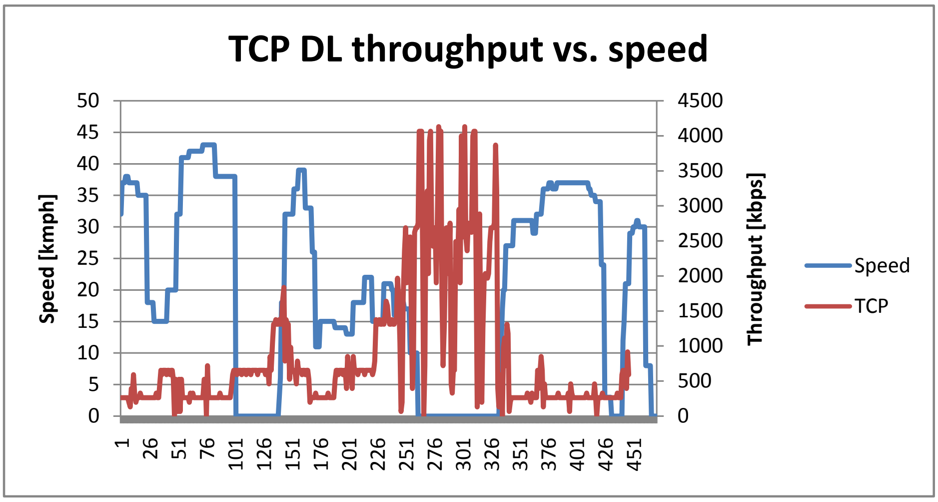

Figure 6.14: TCP DL throughput from BS1 and the speed of the tram

speed of the tram. While in motion, the achievable data speeds seldom rose above 1000 kbps. However, as soon as the tram stopped for a longer period of time, the throughput increased considerably. One example of such behaviour is presented in figure 6.14. It is easily seen that when the tram makes a stop, the throughput increases. As soon as the tram starts to move again, the throughput drops to a very low level. Note that there is some delay from the tram stops until the throughput increases. The reason for this is probably due to TCP and the time it takes to re-build its transmission window after experiencing packet loss, which it falsely believes to be caused by network congestion.

The figure above only represents parts of a single journey with the tram. It is therefore interesting to make a plot of all the measured throughput values as a function of the speed. The result is presented in figure 6.15, which contains data from all the 18 runs, a total of more than 10,000 measurements. This graph clearly states that both high throughput and high speed are mutually exclusive. Generally, throughput above 1500 kbps was very seldom experienced while in motion. When standing still, throughput of more than 4000 kbps was often achieved.

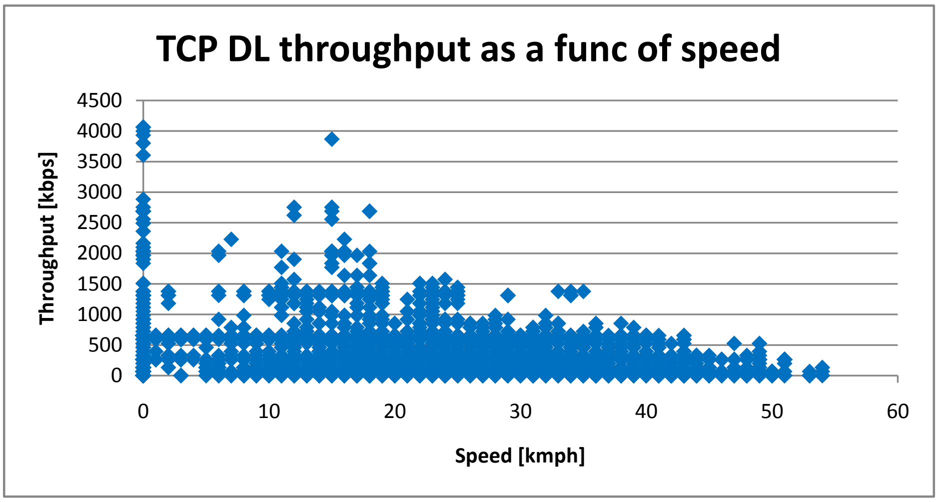

Figure 6.15: TCP DL throughput from BS1 as a function of speed

6.5 BS2: Preliminary Tests Phase I

The next set of tests were performed with the base station installed at site BS2. Initially, no diversity was configured, hence a single vertical polarised antenna was used in the sector. A total of fourteen different locations were visited, where both signal quality measurements and throughput measurements were carried out.

Figure 6.16 shows the locations of the points from where measurements were taken, and table 6.5 provides numerical data. As can be seen, the results were rather depressing. Not from a single point along the last two kilometres of the track did the MS manage to connect to the BS. At points 8 and 9, which both lie in the main lobe of the antenna at a distance of only one kilometre from the base station, the signal was too weak to achieve any throughput at all.

Because of the poor results, it was decided to move the base station antennas closer to the edge of the roof, as their current position could cause unwanted signal reflection and degradation. They were moved about 10 meters further west, and no other changes were performed. Next, new measurements were carried out from the same locations to check for any improvements.

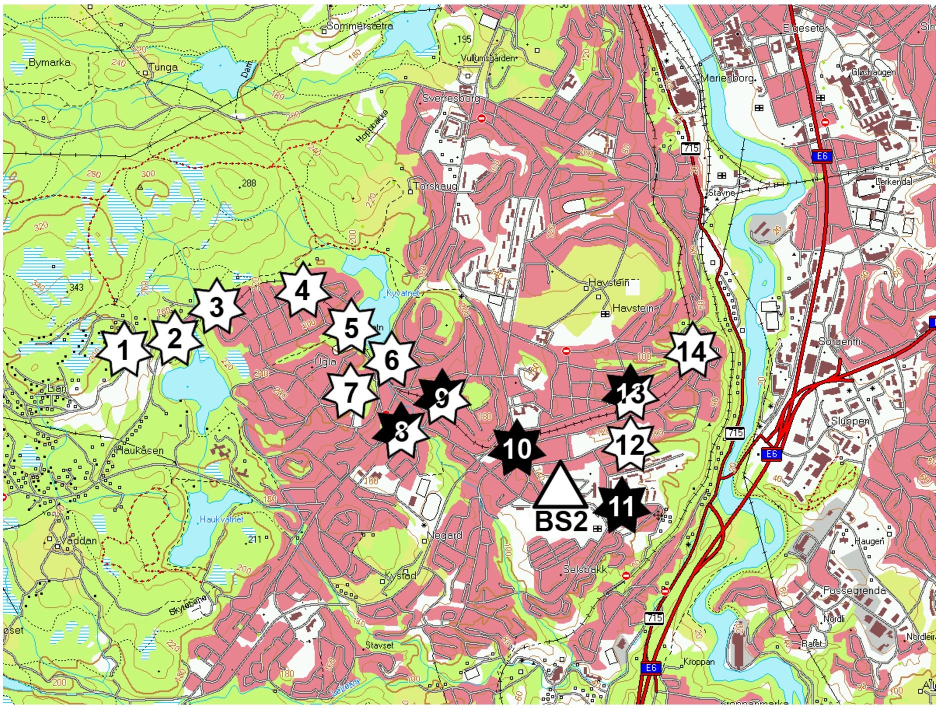

Figure 6.16: Overview of preliminary testing from BS2. White stars indicate no connection. Black stars indicate successful connection. The black and white stars indicate successful connection to the BS, but with a signal too weak for useful communication.

Location	RSSI [dBm]		SNR [dB]		TX Power [dBm]	Distance [km]
	DL	UL	DL	UL		
1						2.80
2						2.70
3						2.50
4						2.00
5						1.60
6						1.40
7						1.40
8	-93	-91	10	10	17	1.10
9	-101	-94	1	6	19	0.97
10	-79	-74	23	26	17	0.36
11	-92	-89	10	12	18	0.37
12						0.48
13	-97	-94	6	7	19	0.66
14						1.20

Table 6.5: Physical performance results from preliminary testing from BS2 Phase I

6.6 BS2: Preliminary Tests Phase II

The new tests showed somewhat better results, as illustrated in figure 6.17 and detailed in table 6.6. This time, the MS was able to successfully connect to the BS from location 1 and 7, while losing connectivity from location 13.

In the following, a comparison of the uplink and downlink SNR values measured during phase I and phase II is presented. Figure 6.18 provides a graphical illustration, where the blue bars represent downlink SNR values and the red bars represent uplink SNR values.

The first observation to make is that the DL SNR values increase considerably from phase I to phase II. This is illustrated with the two different shades of blue, where the light blue bars show the measured results after moving the antenna elements. The average improvement is more than 8 dB, which translates to nearly a three-fold increase in SNR. The UL SNR values does not increase as much, with an average improvement of only 1 dB. This means that the original antenna placement severely distorted the transmission from the base station, while reception was not affected much.

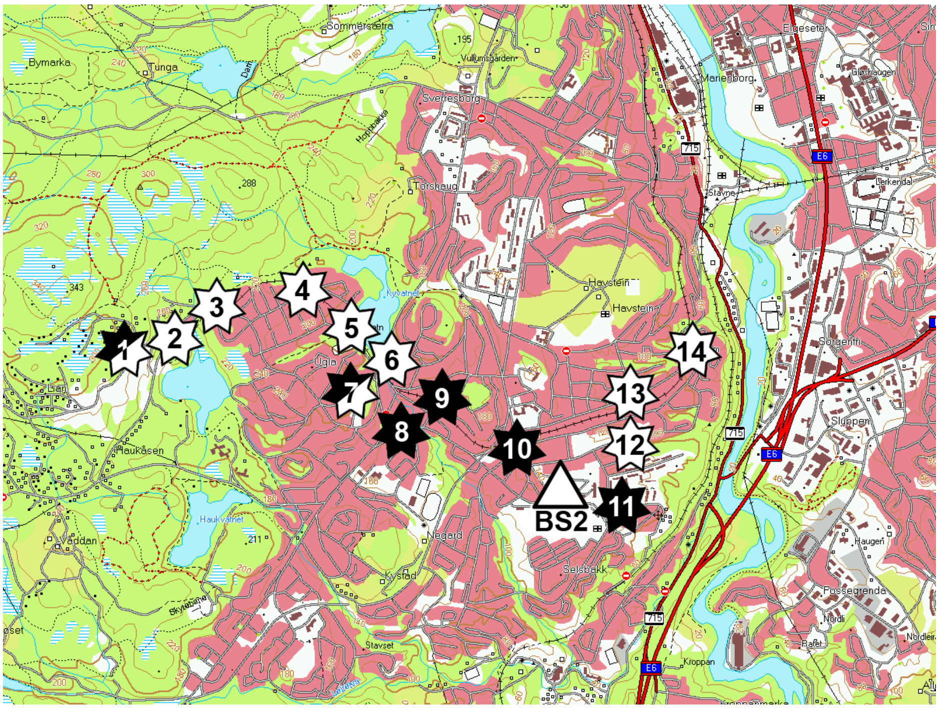

Figure 6.17: Overview of preliminary testing from BS2. White stars indicate no connection. Black stars indicate successful connection. The black and white stars indicate successful connection to the BS, but with a signal too weak for useful communication.

Location	RSSI [dBm] DL	UL	SNR [dB] DL	UL	TX Power [dBm]	Distance [km]
1	-93	-91	9	9	18	2.80
2						2.70
3						2.50
4						2.00
5						1.60
6						1.40
7	-99	-95	4	8	18	1.40
8	-83	-90	19	11	17	1.10
9	-90	-91	12	9	19	0.97
10	-66	-74	31	25	18	0.36
11	-84	-86	16	14	15	0.37
12						0.48
13						0.66
14						1.20

Table 6.6: Physical performance results from preliminary testing from BS2 Phase II

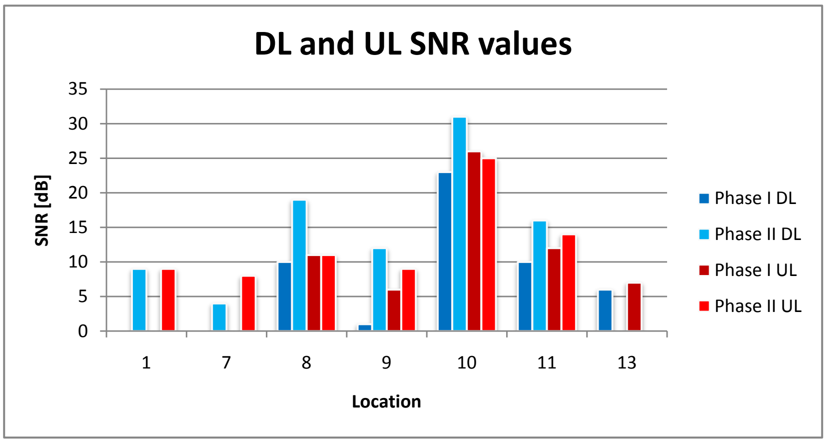

Figure 6.18: Measured SNR values for uplink and downlink during preliminary testing phase I and phase II

6.6.1 Comparison with Propagation Modelling

Comparing the measured coverage with the results from the propagation modelling phase, given in figure 4.8, yields the following.

Location 1 is just within the coverage area, and the measured signal strength here is in accordance with what was expected. At location 2, which lies next to the end station of the tram, but some tens of meters below location 1, no signal was detected even though the propagation model predicted coverage here. Location 3, 4, and 5 all lie outside the expected coverage area, hence not obtaining connection from here is as predicted. According to the propagation model, location 6 lies on the border of the coverage area, hence the lack of signal here is clearly worse than anticipated. The measured downlink signal strength at location 7 was -99 dBm, which is much less than the expected -85 dBm. The cause of this may be dense vegetation between the client and the base station.

At location 8, which lies in the main lobe of the antenna, a signal strength of -76 dBm or better was expected, as opposed to the measured -83 dBm. Location 9 and 10 performed as predicted by the propagation model. At location 11, which lies in the back-lobe of the antenna, no coverage was predicted. However, the downlink signal was measured at -84 dBm, which indicates that signal propagation by reflection in this area is much better than anticipated. Location 12, 13, and 14 are all outside of the coverage area, hence not obtaining connection from here is in accordance with the modelling.

Based on the measured values, the conclusion is that the modelled coverage area to a large degree represents the real world. Some deviation was found, where the measured signal strength was much less than what was expected. The same reasons as given in section 6.3.1 also applies here. In addition, the following causes may account for some of the degradation:

- The base station antennas were not optimally installed with regards to reflections from the roof, even after moving them to a new position.
- The buildings in front of the site, as shown in figure 5.7 blocked the signals more than expected.
- Vegetation was more dense and the landscape more sub-urban in this area compared to the first part of the track. This caused signal propagation to suffer.

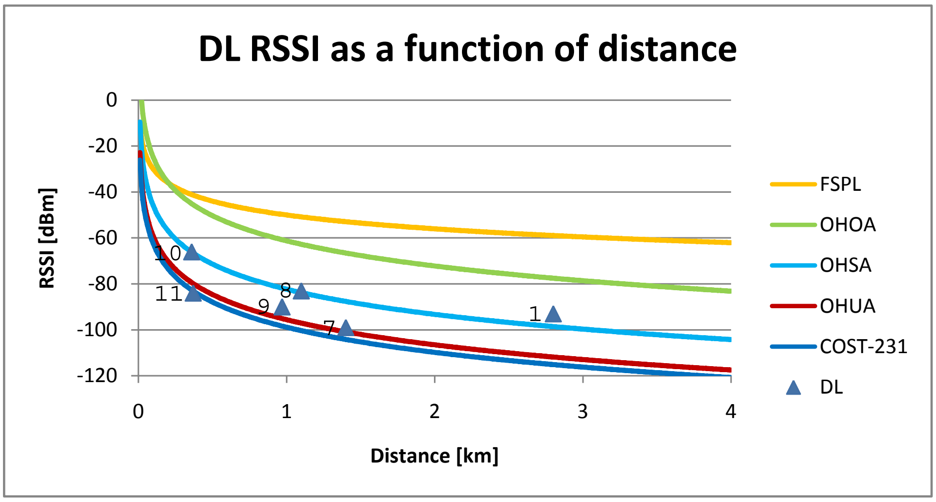

Figure 6.19: DL RSSI as a function of distance from BS1 compared to path loss models

6.6.2 RSSI as a Function of Distance

Figure 6.19 shows the measured DL RSSI values plotted as a function of distance to the base station. The five path loss models are also shown for comparison, calculated using the simple link budget derived in equation (6.1).

It is interesting to note that the measurements taken at location 1, 8, and 10 are according to the Okumura-Hata model for sub-urban areas, while the three others are according to the urban model. This may be because both point 7 and 9 reside in areas covered with dense vegetation, and point 11 is located in the back-lobe of the antenna without a clear line of sight to the base station. Because of this, the urban model makes more sense for these locations.

6.6.3 SNR as a Function of RSSI

The plot of the SNR as a function of RSSI is shown in figure 6.20. The linear relationship between the two parameters is clearly seen, where an increase in SNR yields a similar increase in RSSI. Because all the measurements fall close to the regression line it can be stated that the noise in the channel is relatively constant.

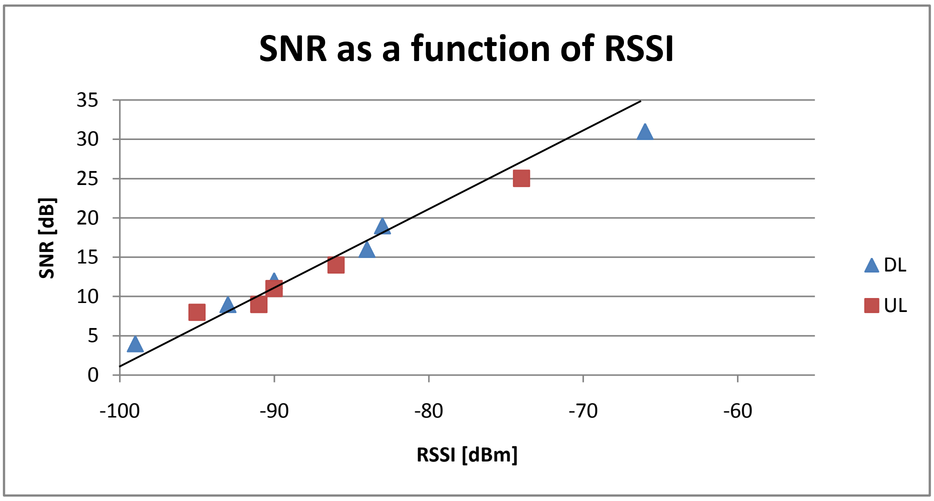

Figure 6.20: SNR as a function of RSSI for both uplink and downlink

Location	DL throughput [Kbps]	UL throughput [Kbps]
1	310	N/A
7	284	N/A
8	851	N/A
9	395	N/A
10	3062	N/A
11	333	N/A

Table 6.7: Average TCP throughput from preliminary testing from BS2

6.6.4 TCP Throughput as a Function of RSSI

The TCP DL throughput was measured at all locations from where a connection to the base station was successfully achieved. The test was run three times to mitigate the effects of intermittent errors. TCP UL throughput was not measured. The averaged numerical results are presented in table 6.7 and graphically illustrated in figure 6.21.

Compared to earlier tests, the achieved throughput is similar for identical values of RSSI. However, the data throughput is still much less than what is theoretically possible. For instance, at location 8 the measured DL and UL RSSI were equal to -83 dBm and -90 dBm respectively. In the downlink, a modulation rate of QAM 3/4 should be possible, yielding a maximum gross throughput of 12096 kbps. As the UL/DL ratio is set to 50/50, only half of this data rate is possible. Still, only 851/6048 $\approx$ 14% efficiency is achieved in the downlink.

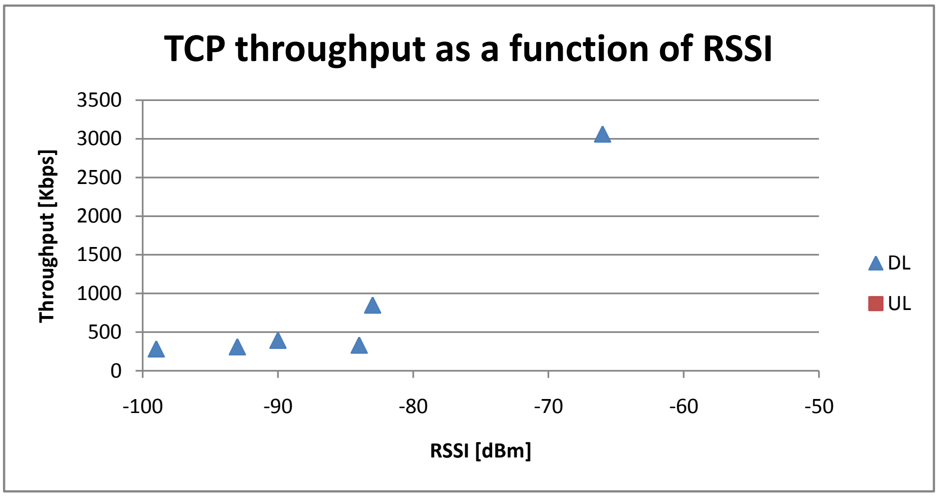

Figure 6.21: TCP throughput as a function of RSSI

The same thing goes for location 9. The measured DL and UL RSSI are equal to -90 dBm and -91 dBm respectively. A throughput of 395 kbps is achieved, much less than the 2 Mbps theoretically possible.

6.7 BS2: Testing on the Tram

The last set of tests were carried out on the tram. Like before, several journeys were made from start to end while performing measurements on the last part of the track. Testing was carried out both without diversity and with 2^{nd} order diversity using separation in space. Separation by polarisation was not tested. The results from the testing are presented in the following sub-sections.

6.7.1 Signal Reach

The points where the signal was picked up and lost have been plotted on a map to verify the coverage area and to check if 2^{nd} order diversity caused any improvements. Several journeys were made in both directions to make sure that the measurements taken were realistic.

Figure 6.22 shows the results from Hoem. The red and blue diamonds indicate where the signal was picked up on the trips upwards using no diversity and 2^{nd} order diversity respectively. It can easily be seen that diversity has

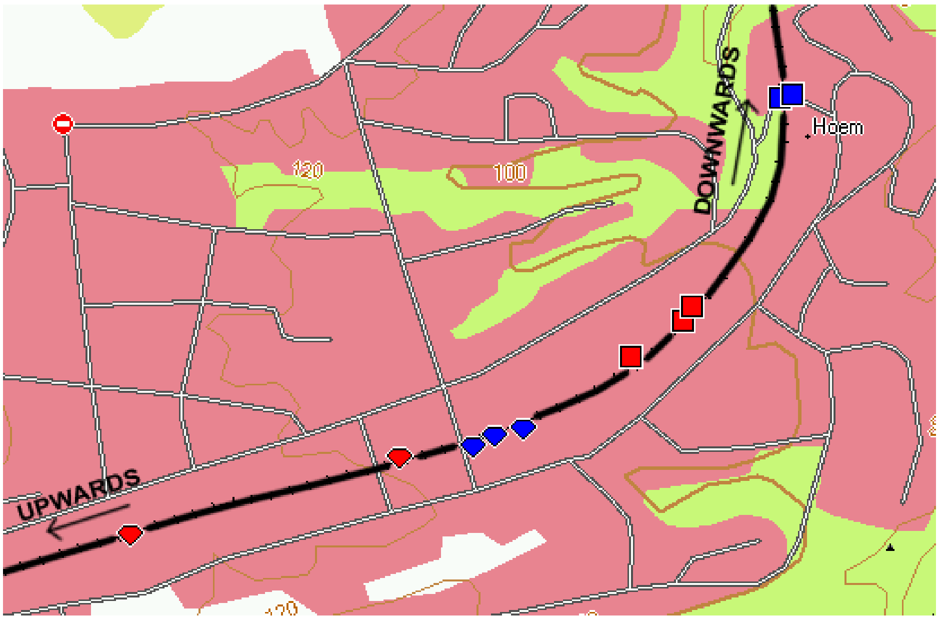

Figure 6.22: Signal reach at Hoem. Diamonds indicate where the signal was picked up on trips upwards, squares where it was lost on trips downwards. Red means no diversity, blue means 2nd order diversity using VV.

some impact on the reception, as the signal is picked up earlier. Also, while the blue diamonds lie close together, the red diamonds are more spread out. When it comes to losing the signal, diversity leads to an increase of about 300 meters. It is interesting to note that each group of squares lie close together.

Comparing the measured results to the propagation modelling in figure 4.8, it can be seen that the signal reach is much better than predicted by the model. This is somewhat contradictory to what was measured during the preliminary testing, as none of the locations visited in this area obtained connectivity.

The results from Kyvatn are presented in figure 6.23. Once more it is clear that diversity improves the coverage area, as the signal is picked up earlier and lost later. However, the actual distance of improvement is very small, measuring only 150 meters. Looking at figure 4.8, the measured signal reach in this area coincides very well with what was predicted by the propagation model.

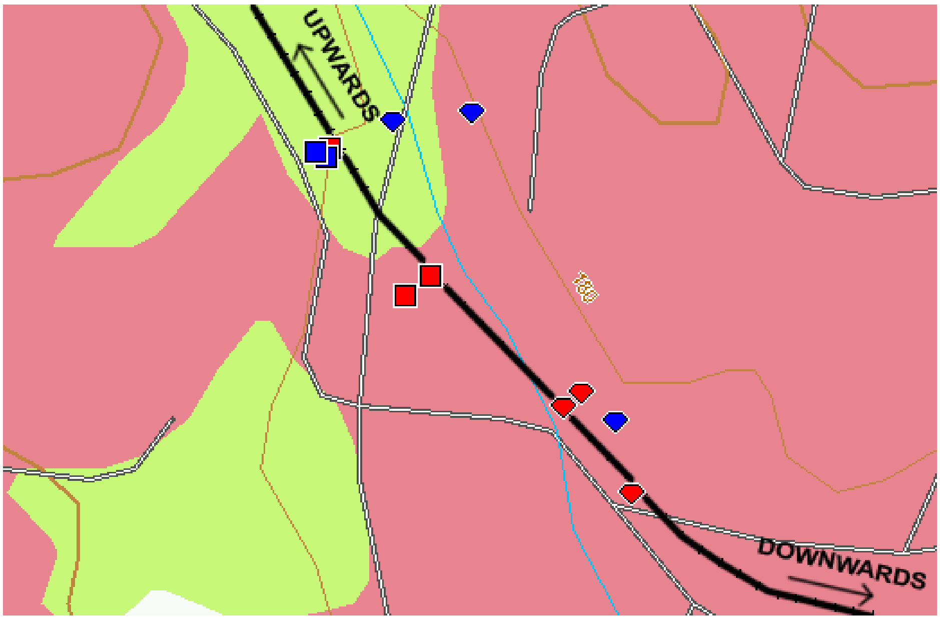

Figure 6.23: Signal reach at Kyvatn. Diamonds indicate where the signal was picked up on trips downwards, squares where it was lost on trips upwards. Red means no diversity, blue means 2nd order diversity using VV.

Diversity	DL RSSI [dBm]		UL RSSI [dBm]	
	Upwards	Downwards	Upwards	Downwards
None	-83	-83	-85	-85
VV	-77	-82	-81	-84

Table 6.8: Average physical performance measurements from BS2

6.7.2 RSSI as a Function of Distance

The average DL RSSI results from 12 journeys on the tram are presented in figure 6.24. As in the previous tests, the values have been averaged in order to mitigate intermittent errors. The blue line represents the results from using no diversity, while the red line represents 2nd order diversity using separation in space. Separation by polarisation was not tested from BS2. Numerical values are presented in table 6.8, separated by direction upwards and downwards.

According to the graph, there is a clear indication that 2nd order diversity leads to some improvements, as the red line generally lies above the blue line most of the time. This is confirmed by looking at the numerical values in the table, which states that higher values of both DL and UL RSSI are achieved

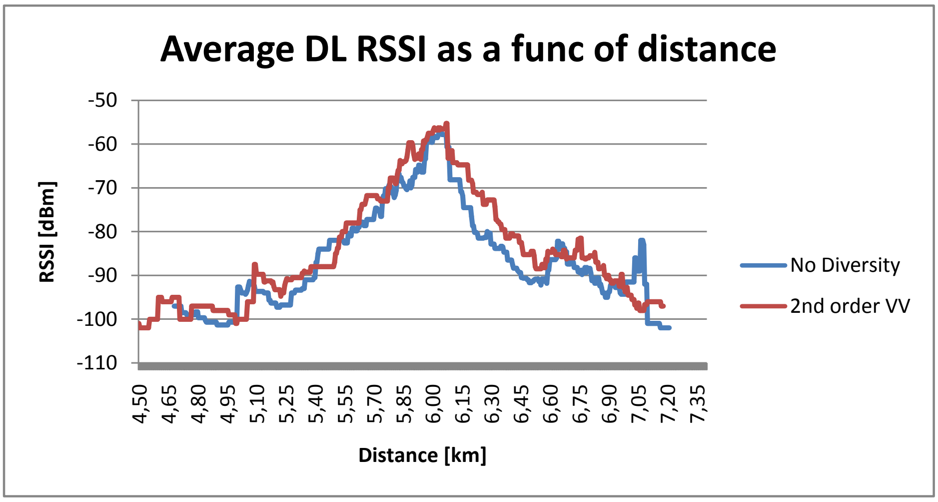

Figure 6.24: Average DL RSSI values from 12 consecutive runs with and without diversity from BS2

when diversity was used. Without diversity, there is a difference of 2 dB between DL and UL RSSI, where the downlink achieves the higher value. The user equipment may be accountable for this, as the radio transmitter in the CPE is not as powerful as the one located at the base station. It is also interesting to note that there is no difference in the direction of travel, as both upwards and downwards journeys achieve identical values.

In the case of diversity, there is much more difference. While the journeys downwards only achieve a 1 dB increase in RSSI, upward journeys report an increase of 6 dB and 4 dB for DL and UL respectively. As commented on earlier, the antenna placement on the tram may be accountable for this difference.

The strongest values of DL RSSI were measured at approximately 6 km along the track. Referring to figure 4.1, this area lies right in front of the antenna lobe, with a distance of about half a kilometre from the base station. Values of -60 dBm were measured in the downlink, which translates to a very strong signal. However, the RSSI falls rapidly on either side of the 6 km mark.

The graph also states that the signal is picked up earlier when diversity is used. This fact was confirmed in the previous sub-section dealing with signal reach. At the other end, the two configurations perform just as well, which is in accordance with what was expected.

Before ending the discussion on RSSI, a plot of the DL RSSI values as a

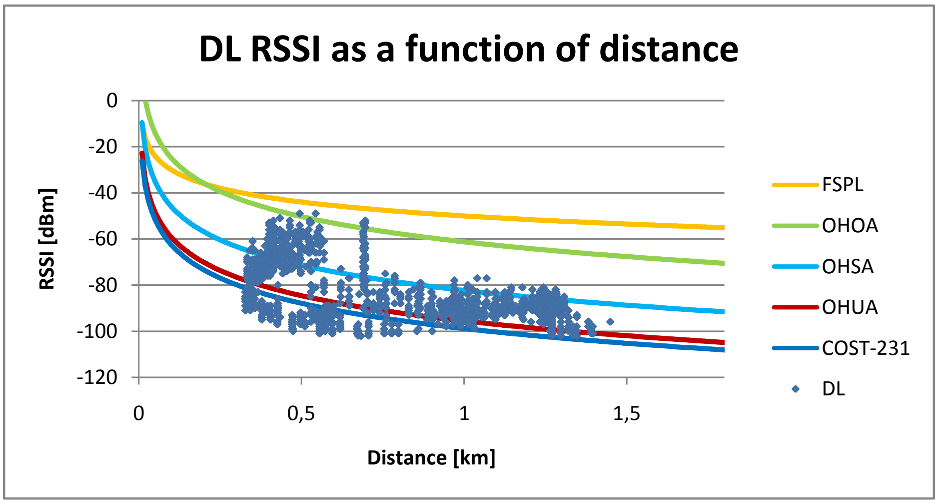

Figure 6.25: DL RSSI as a function of distance from BS2 compared to path loss models

function of the distance between the tram and the base station is presented in figure 6.25. The five different path loss model are also included for comparison. The plot is based on six journeys along the tram line, three in each direction. No diversity was used, and a total of more than 2000 measurements were taken. There are several interesting points to make from studying these results. First, the DL RSSI values are generally much lower here compared to testing from BS1. Second, the distance to the points farthest away from the base station is less than 1.5 kilometre, which makes it only half of what was achieved on the first part of the track. Third, except for the relatively strong signal strength measured at half a kilometres from the base station, most measurements lie in between or below the Okumura-Hata models for sub-urban and urban access.

6.7.3 TCP Throughput

The throughput of a TCP DL connection was also measured both with and without diversity configured. The results are graphically illustrated in figure 6.26, and table 6.9 contains the numerical data separated by both diversity order and direction of travel. The results are based on three round-trip journeys using no diversity and the same number using 2^{nd} order diversity, yielding a total of 12 runs.

The results from the testing show that the throughput increases slowly from

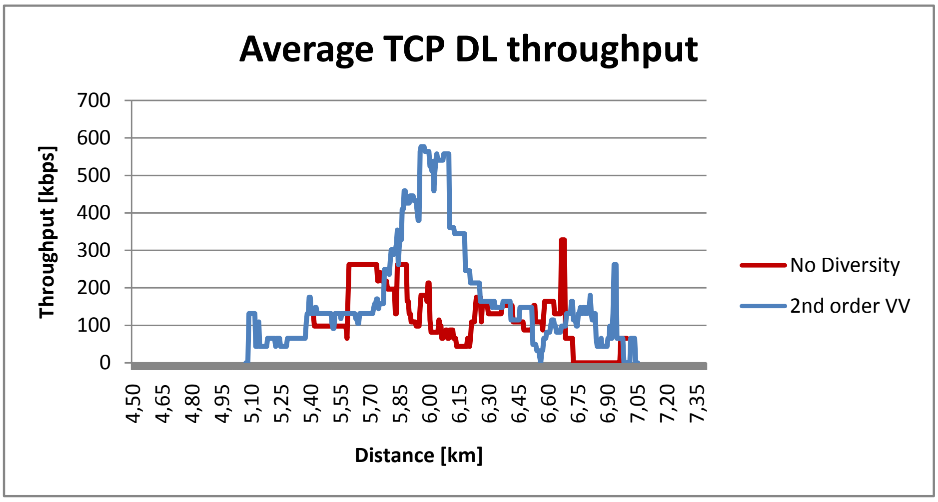

Figure 6.26: Average TCP DL throughput from BS2 as a function of distance along the tram line. The red line illustrates no diversity, and the blue line represents 2nd order diversity using separation in space

Diversity	TCP Throughput [kbps]		
	Upwards	Downwards	Average
None	94	166	130
VV	211	184	198

Table 6.9: Average TCP throughput measurements from BS2

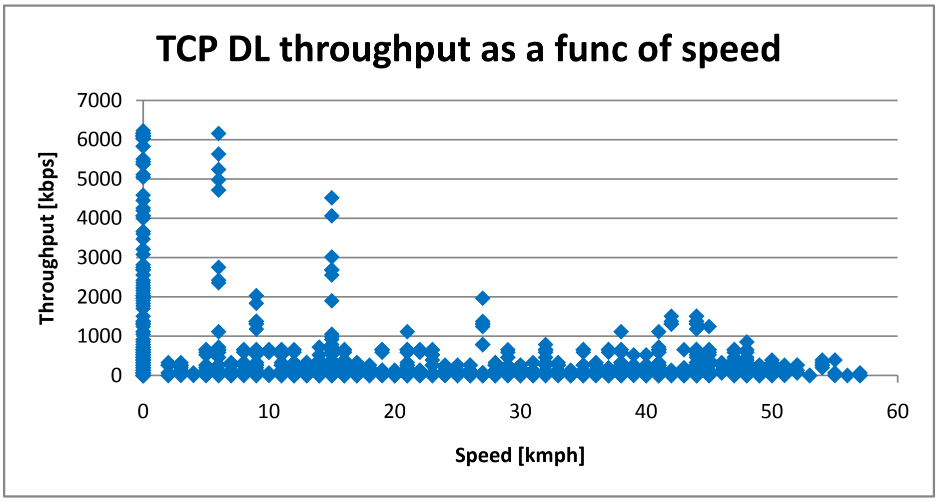

Figure 6.27: TCP DL throughput from BS2 as a function of speed

the start, averaging between 100 and 200 kbps. At about 6 km, the throughput reaches a maximum with average values between 400 and 600 kps for second order diversity. When using no diversity, there is no identifiable maximum. Referring to figure 6.24, which plots the average DL RSSI values, 6 km is also the area where the signal strength is highest. Onwards from here, the throughput steadily decreases, with the exception of a small increase at 6.6 km.

Generally, the throughput achieved from using no diversity is much less than for 2$^{\text{nd}}$ order diversity. This may have several causes, but prime among them is the signal strength. The measured DL RSSI was significantly stronger when diversity was in use. Hence, a strong signal leads to a higher order of modulation, and thus more bits may be transferred per symbol.

Also note that there is a clear difference between where the signal was picked up and where data started to flow through the system with respect to distance. Comparing figure 6.24 to figure 6.26 yields that even though the CPE connected to the base station at approximately 4.5 kilometres when using diversity, no throughput was measured until 5.10 kilometres, a difference of 600 metres. The same goes for the runs without diversity, where the difference in length is a bit more.

Earlier in this thesis it was observed that the speed of the tram and the achievable throughput were highly correlated. It was noted that the throughput increased noticeably when the tram stopped, while it decreased as soon as the tram started moving again. Identical observations were made during

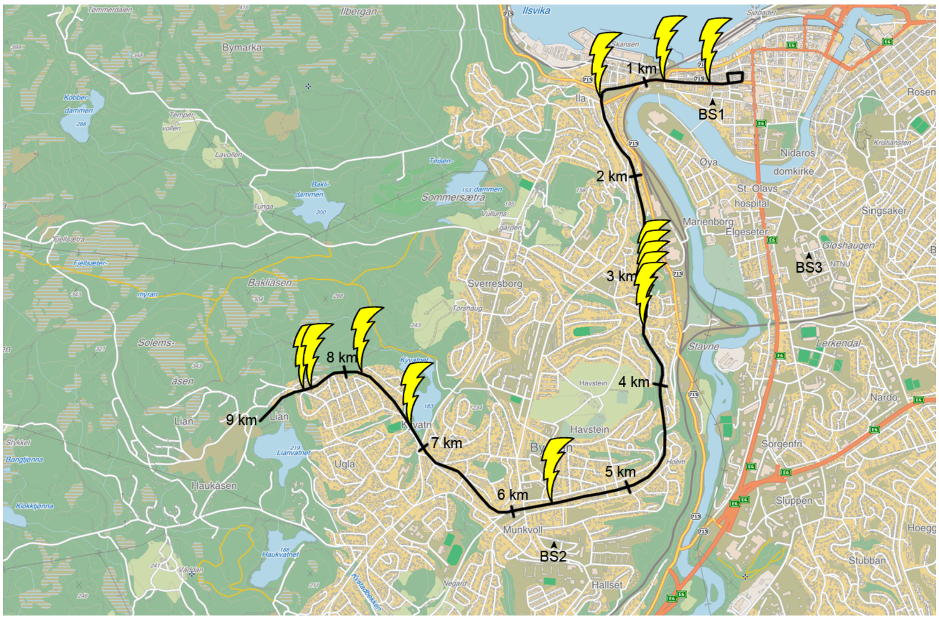

Figure 6.28: Map of the tram line illustrating points where loss of power occurred

testing from BS2, and a plot of the throughput as a function of speed is presented in figure 6.27. It can clearly be seen that throughputs of more than 6000 kbps were archived while waiting at a stop. However, throughputs of more than 1000 kbps were seldom experienced while in motion.

6.8 Problems Encountered

One of the more frustrating problems that occurred was sudden loss of power in the tram. The laptop computer handled the loss well, as it immediately switched to battery power. But as the CPE required constant and stable power supply, and did not have a battery for backup, power outages caused the CPE to shut down and reset. Because it took about a minute for the client equipment to boot and re-connect to the base station, during some runs, as much as a minute of measurement data were lost for each outage. These periods can easily be spotted as blank areas in the raw testing material. However, as the graphs presented in this thesis use averaged values, the gaps are not visible.

A map of the tram line indicating the places where power loss occurred is given in figure 6.28. As can be seen, a lot of outages happened at the 3 km mark, while the others were spread over the entire line.

There are three main reasons for such loss in power. First, at the 3 km mark, and between the 5 and 6 km marks, there are section breaks where a small part of the overhead line is insulated and thus disconnected from the supplies. If the tram crosses these parts with low enough speed, the time spent in the electrically dead area is sufficient to shut down the CPE. Second, when two trams, both a long distance away from the feeder stations, simultaneously draw a lot of power, the voltage drops to a level where the user equipment cannot function, and hence shuts down. According to the tram operator, this phenomenon is particularly evident in the inner city. Third, the pantograph may occasionally loose contact with the overhead lines, which may happen at crossings.

Chapter 7

Discussion and Recommendations

In this chapter, a general discussion with regards to the results is provided.

7.1 Coverage Without Diversity

With respect to coverage without diversity in use, the propagation model was found to represent the real world quite well. Especially testing with the base station installed at site BS1 showed very promising results, with measurements lining up very closely to what was predicted by the model. Some deviation was found, where the measured results were a bit less than what was anticipated. Especially along the first kilometre of the track, which ran through the inner city. The reasons for not achieving optimal results here were mostly due to the urban environment, where tall buildings blocked signal propagation, and the fact that this area was located in the back- and side-lobes of the antenna. Also, because no building height file was used during modelling, and the clutter layer was a bit coarse, the estimated results were a bit optimistic.

Just like predicted by the model, the sector from BS1 did not provide adequate coverage to the starting point of the tram line. Also, the signal quality from the 3 km mark and beyond varied much, with average DL RSSI values between -80 dBm and -90 dBm. This was illustrated in figure 6.11. Because a signal strength of more than -80 dBm is required in order to provide sufficient throughput, it was concluded that the sector from BS1 only provided adequate coverage of the tram line between the 1 km and 3 km marks. If the required signal strength is lowered to -90 dBm, the entire track from the

0.5 km mark until the 4.5 km mark may be covered by this sector, with the exception of some loss at 3.6 km, where the tram line makes a turn around a small hill.

It was also found that the DL RSSI measurements taken along this part of the track mostly lined up in between the Okumura-Hata models for open and sub-urban areas. This was illustrated in figure 6.12. Requiring a signal strength of at least -80 dBm yields a maximum distance to the base station of 1.5 km. This corresponds to what was found above, where the 3 km mark along the tram line is exactly 1.5 km away from the base station. Lowering the requirements to -90 dBm yields a maximum distance of 3 kilometres, corresponding to the 4.5 km mark.

Testing from site BS2 yielded results that were a bit less than what was predicted by the model. Especially measurements taken in front of the main lobe reported DL RSSI values much lower than expected. Possible reasons for this may be related to the buildings in front of the antenna which blocked signal propagation, and relatively dense vegetation in the area. The placement of the base station antennas was also commented upon. When it comes to coverage of the tram line, the area predicted by the model lines up very well with reality. Assuming that a signal strength of more than -80 dBm is necessary to provide sufficient throughput, the sector only covered the area between the 5.5 km and 6.3 km marks. In other words, a relatively small area of less than one kilometre. This was illustrated in figure 6.24. If the required signal strength is lowered to -90 dBm, the track length between the 5.3 and 6.8 km marks may be covered, almost doubling the length.

The DL RSSI measurements taken along this part of the track was found to line up in between the Okumura-Hata models for sub-urban and urban areas. This was illustrated in figure 6.25, and it clearly showed that the signal strengths measured here were much lower than for the first part of the track. Requiring a signal strength of at least -80 dBm yields a maximum distance to the base station of approximately half a kilometre. In other words, the sector from BS2 did not provide very good coverage of the tram line.

Comparing the results obtained here to the related work presented in the beginning of this thesis yields the following. In [1], a self-installable indoor unit with a 7 dBi antenna was used in the 3.5 GHz band. The maximum distance the unit was able to operate under was 1.5 kilometres from the base station. As both the testing environment and the antenna gain are comparable to this project, it can be seen that NLOS coverage of more than 1.5 kilometres from the base station in sub-urban environments seem unlikely. Comparable results from [2] state that full throughput was achieved only within 1 kilometre from the base station using the very same client inside a city.

The testing in [1] also concluded that differences of almost 10 dB were measured in signal reception between spring and summer, mostly due to the increased level of foliage. As the testing in this project was carried out during early summer, and the entire area was covered with dense trees, it is estimated that better results may be achieved during winter.

7.2 Diversity Improvements

When it comes to picking up the signal, 2^{nd} order diversity using separation by polarisation showed the most promising results. Separation in space provided more fluctuating outcomes, while no diversity clearly underperformed. However, the actual differences between the three configurations were rather small, measuring only a few hundred meters. Loss of the signal happened at the exact same spot both with and without diversity in use, hence no improvement was noted here.

The measured RSSI values for both uplink and downlink showed varying results. Measurements taken with the base station installed at site BS1 using diversity in space showed a noticeable increase of 4 dB in DL RSSI for trips upwards. For trips downwards, regardless of diversity order, almost identical results with little variation were noted. The measured results for UL RSSI showed that diversity in fact produced lower or identical values, hence no enhancement was found here.

Measurements taken with the base station installed at site BS2 were more in line with theory. A noticeable increase of 6 dB and 4 dB in RSSI were measured for DL and UL respectively on trips upwards. On trips downwards, the difference was only 1 dB in favour of diversity.

Looking at the outcome of the testing as a whole, it is realised that 2^{nd} order diversity increased the received signal strength when the tram moved upwards, while little difference was measured on trips downwards. As mentioned earlier in this thesis, the antenna placement on the tram may be accountable for this difference. It may be concluded that diversity did only have a small impact on coverage, as the results obtained from testing varied much.

7.3 Throughput

The overall conclusion from testing is that the throughput measured was much less than expected. With the current settings of using a 5 MHz channel and an UL/DL ratio of 50/50, throughput of up to 9 Mbps should be theoretically possible, while only 6 Mbps was achieved when standing still. In motion, throughput seldom rose above 1 Mbps. There are many possible explanations, and they will be discussed in the following.

First, TCP may not be the protocol best suited for this kind of testing, as it takes a bit of time to rebuild speed after experiencing packet loss. And packet loss is experienced all the time in a wireless environment. However, as was discussed earlier, the user experience is better modelled using this protocol. The send and receive windows of the protocol may also impact the throughput if they are not adjusted to accommodate the BDP. During testing, it was made sure that the windows at each side of the communication link used window sizes much greater than the BDP.

Another possible reason to the limited throughput may be due to protocol overhead. In addition the raw user data, each protocol layer adds to the redundancy by including an additional set of headers which must be transported over the wireless link. This causes the effective throughput of the link to drop below the gross throughput. However, as discussed in section 2.6, the IEEE 802.16 standard specifies a method known as packet header suppression, which removes the repetitive part of the protocol headers in order to enhance the data throughput rate. It is not clear whether or not the WiMAX system used during testing implemented this functionality, but the method is listed as mandatory in the mobile WiMAX system profile. If PHS was not implemented, protocol overhead may be accountable for parts of the low throughput.

A third possible reason may be the transfer of management frames between the CPE and the base station. However, it is not likely that the amount of management frames transferred had such an impact on the throughput as experienced here.

All the aforementioned causes may be related to the limited throughput experienced while standing still. In motion, the throughput degradation was so severe that it requires other explanations. In general, the physical throughput depends on the SNR, the BER, and the bandwidth. Their relationship is described in figure 7.1 and explained next.

An increase in distance between the client and the base station generally

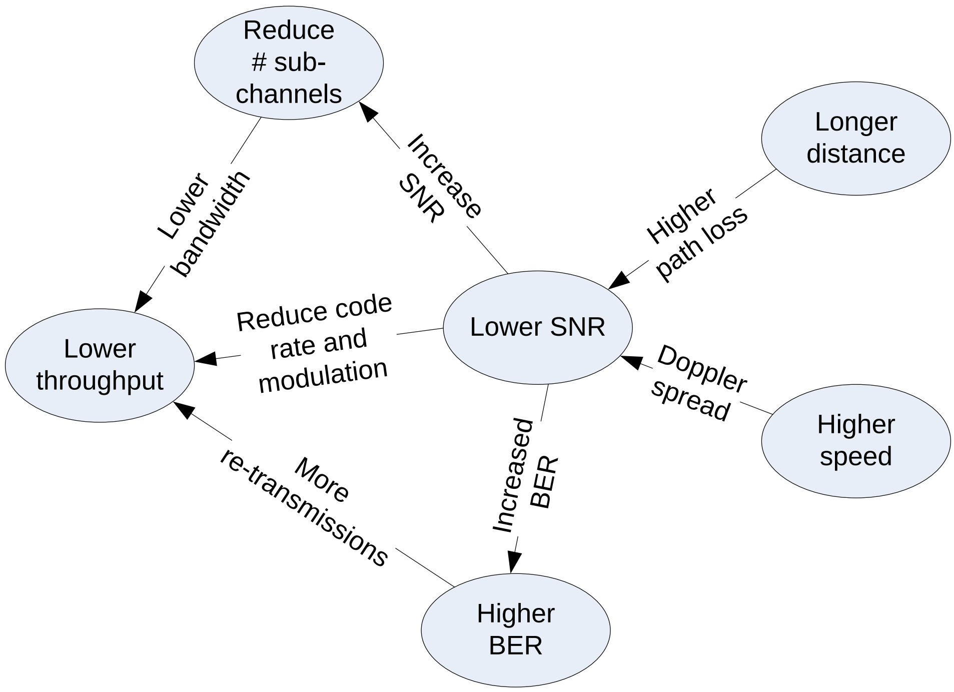

Figure 7.1: Causes of low throughput

leads to higher path loss, which causes lower SNR. This loss may be compensated for in different ways. For instance, the client may choose to reduce the number of sub-channels used in the uplink in order to increase its transmission power. This would increase the UL SNR but at the same time lower the bandwidth. The result is that the throughput is sacrificed for higher availability. Note that sub-channelisation is not supported by the system in the downlink, only in the uplink. In other words, this is only an option for the client, not the base station.

Another alternative is to keep transmitting on all the sub-channels, but reduce the code rate and order of modulation. This reduces the throughput as less bits are transmitted per symbol. However, the availability is increased as the receiver more easily interprets the symbols correctly.

If the number of sub-channels cannot be reduced, which is the case in the downlink or if the number of sub-channels used by the client in the uplink is already one, or the code rate or modulation cannot be changed to a lower order, which is the case if BPSK $1/2$ is already being used, the BER increases. The result is erroneous frames and more re-transmissions. As a consequence, the throughput falls to a very low level.

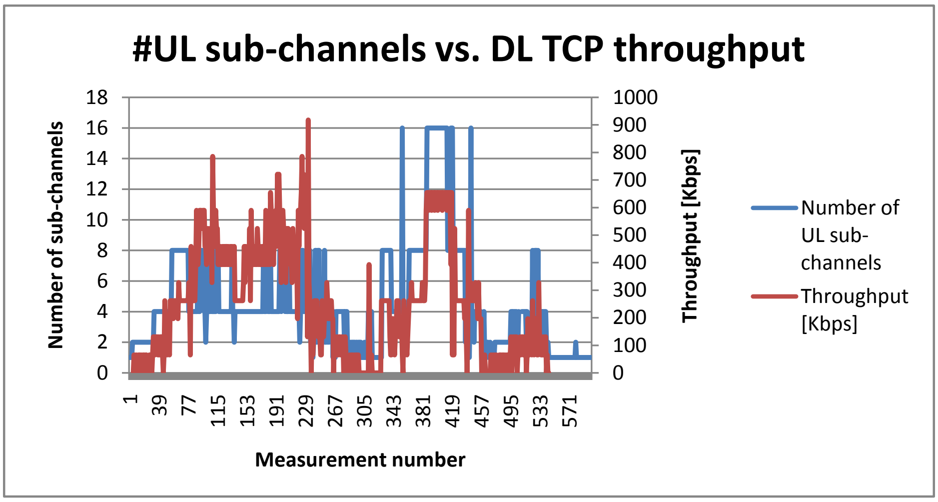

Figure 7.2: DL TCP throughput versus the number of UL sub-channels in use by the client

Now, what about speeds? It was noted earlier that the throughput dropped noticeably when the tram was in motion, and that only when the tram stopped were data rates above 1 Mbps achieved. In the following, one particular part of the tram line is analysed with regards to throughput and speed, where measurements are taken every second. All the plots cover the same area, hence they may be compared against each other.

A plot of the downlink TCP throughput and the number of sub-channels in use is presented in figure 7.2. Here, the blue line illustrates the number of sub-channels while the red line illustrates the throughput. As can be seen, there is a high correlation between them. This is quite interesting, as the throughput is measured for the downlink, while the sub-channels represent traffic going in the opposite direction. This suggests that TCP relies on the link in both directions. It is also interesting to note that the number of sub-channels vary widely. When few are in use, the throughput is low, and vice versa.

The next plot, which is given in figure 7.3, illustrates the number of sub-channels and the speed of the tram. As was noted above, the number of sub-channels vary with time. It seems that there is a relatively strong correlation between the speed of the tram and the number of sub-channels in use. When the speed is zero, a high number is used, while the number drops when the tram starts to move.

Last, there is a plot of the UL and DL SNR values, given in figure 7.4.

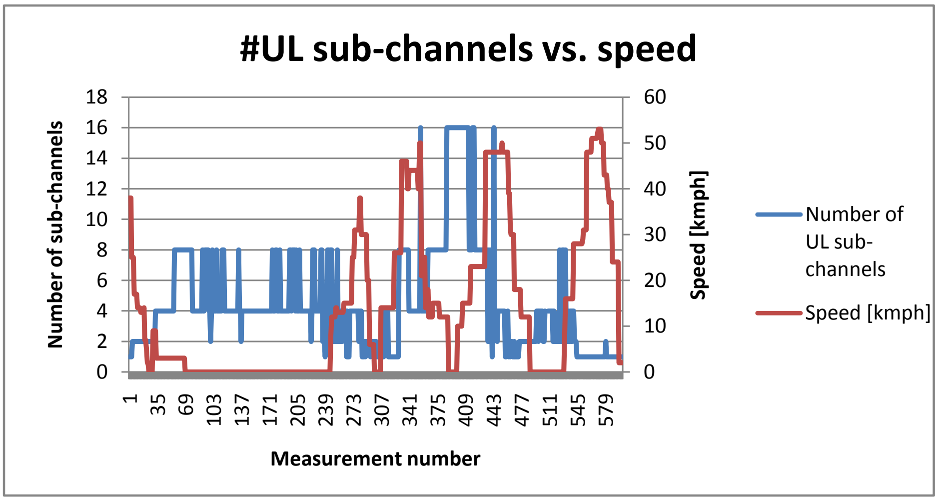

Figure 7.3: The number of UL sub-channels used by the client versus the speed of the tram

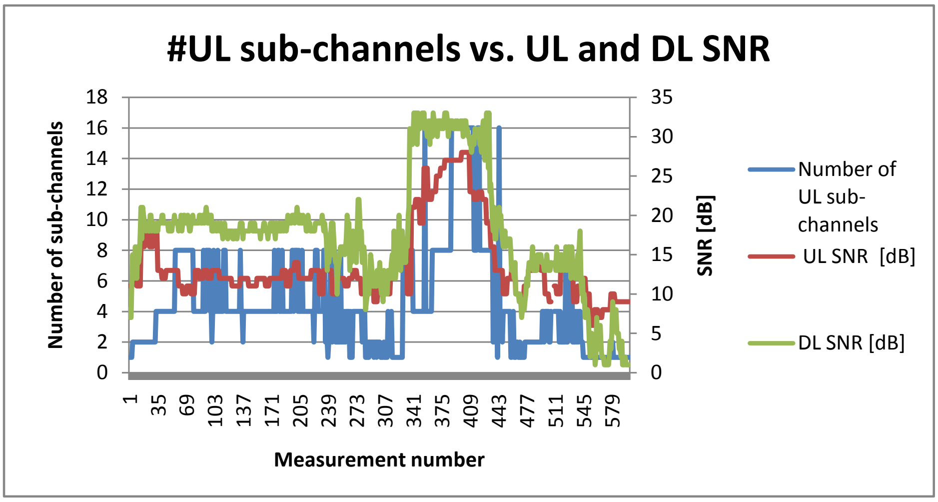

Figure 7.4: Number of UL sub-channels used by the client versus UL and DL SNR

Looking at the red line, which represents the UL SNR, it can be seen that it stays around 12 dB for the first 300 measurements. When the tram starts to move, there is little variation in the reading. However, the number of sub-channels is reduced from 8 to 4, and then again to 2. Consulting figure 7.2 again, it can be seen that the throughput drops from approximately 400 to 200 kbps, and then again to 100 kbps.

As we know, when the distance between the transmitter and receiver varies, Doppler shift is experienced. This causes channel fading, and the higher the speed the more severe the fading. As discussed in [44, 45], OFDM systems like WiMAX will suffer due to Doppler spread because the time variations of a fading channel lead to loss of sub-channel orthogonality, resulting in ICI. If not compensated for, an error floor will increase with the Doppler frequency, dependent on the number of sub-channels and the fading rate.

Assuming that the number of sub-channels in the uplink is related to the TCP throughput in the downlink, the conclusion so far is that when the tram is in motion, the Doppler spread causes ICI and loss of orthogonality is experienced. To compensate for this, the client chooses to utilise less sub-channels, and thus the throughput drops as TCP is dependent on the link in both directions. This subject clearly needs more investigation.

It is shown in [45] that antenna diversity can be used as an effective means of reducing the degrading effects of a fading channel. The less complex methods, such as selection combining, offer no improvement, because ICI is not a function of signal power, but rather the rate at which the channel varies. However, MRRC provides acceptable performance over a wide range of Doppler frequencies. It was observed during testing that 2^{nd} order diversity using MRRC in the uplink led to much higher throughputs. Using separation in space, improvements of 42% and 52% were measured from site BS1 and site BS2 respectively. One possible explanation of this may be that antenna diversity led to improvements with regards to Doppler.

Another possible cause of degradation while in motion may be related to the use of an omni-directional antenna on the tram. According to [46], the use of directional antennas in a mobile channel enables enhanced signal reception, reduced Doppler spread, reduced delay spread, and reduced CCI. It is further shown via computer simulations that the beamwidth and orientation of a directional antenna can limit the Doppler spread, and hence the channel fading rate, as opposed to using an omni-directional antenna.

A field trail was also implemented in [47], where a fixed WiMAX system operating in the 5.8 GHz frequency band was set up in a sea port environment.

Measurements were carried out using vessels with relatively low mobility. It was found that Doppler shift was the main factor degrading the performance, and that a directive antenna mitigated this effect to a large degree.

To conclude this section, it seems clear that the throughput is dependent on at least two factors. The distance to the base station and the speed of the tram. The low throughput that was experienced while in motion seems to be related to Doppler spread. It was shown that diversity somewhat improved on this, as higher throughputs were achieved when MRRC was used. The omni-directional antenna used on the tram seems to have led to some degradation. According to several sources, a directional antenna would have been better.

7.4 Recommendations

The following section provides recommendations for improving the coverage and throughput along the tram line.

7.4.1 Base Station Locations

The sector from the base station at site BS1 seems to provide good coverage of the first few kilometres. No changes are proposed to this sector. In addition, it is recommended to implement another sector at this site pointing eastwards. This would provide sufficient coverage to the starting point of the tram line, as well as to the remaining parts of the inner city.

The sector from the base station at site BS2 does not seem to provide adequate coverage. Also, there is little or no coverage between the 3 and 5.5 km marks, and from the 6.5 km marks and onwards. It is recommended to move the base station to a new site located closer to 5 km mark, and to set up a new sector that could provide coverage to the area north of here. Additionally, a new site would be necessary to cover the remaining parts of the tram line. It should be located as high as possible, in order to provide adequate coverage to the line behind the hill Ugla. One possible location is at the Telenor mast at Storhaug, a solution which was discussed in section 4.3.6.

A sketch of the recommended sectors is given in figure 7.5. This would require two new base station sites, which would increase cost and provide some challenges with regards to high speed backhaul.

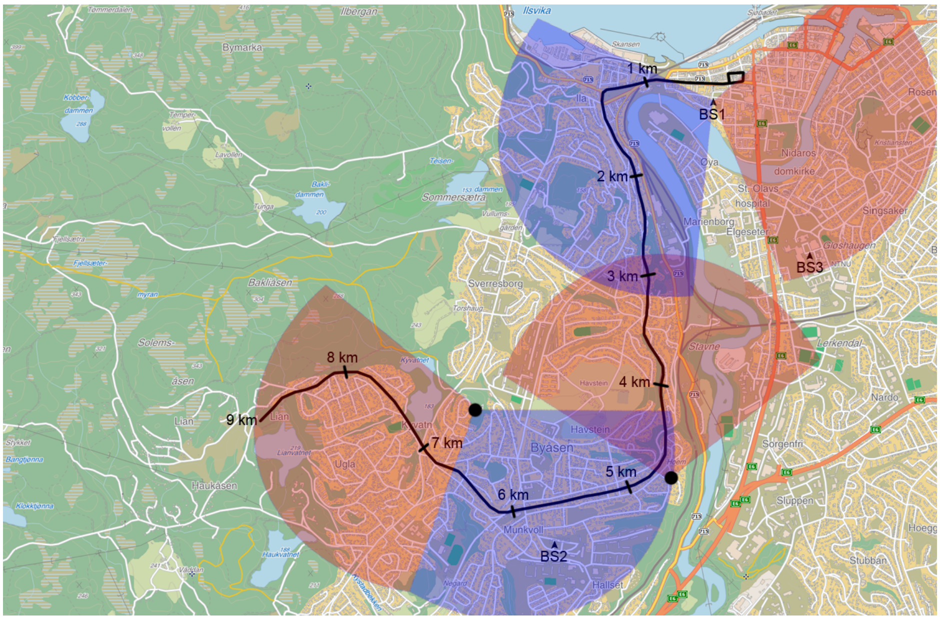

Figure 7.5: Map of the tram line illustrating recommendations for new sectors

7.4.2 CPE Antenna

The antenna used on the tram seems to be sub-optimal for the task. One problem is related to the waste of energy, because of its omni-directional beam pattern. The other is related to degraded service while in motion, because of more sensitivity to Doppler spread. It is recommended that another antenna is implemented on the tram in order to provide better reception and transmission capabilities. For instance, there exists sectorised antenna arrays that provide omni-directional coverage but only utilise the one sector facing the base station when transmitting or receiving.

7.4.3 Power Loss

In order to solve the problems related to power loss on the tram, it is recommended that battery backup is provided that can temporarily feed the CPE when power loss occurs. Because it takes about a minute to restart the equipment, such a loss would severely impact the availability of the service.

Chapter 8

Conclusion and Future Work

This chapter presents the conclusion as well as recommendation for future work in the area.

8.1 Conclusion

It is clear that the IEEE 802.16e-2005 amendment adds valuable improvements to the WiMAX standard. Support for full mobility along with enhanced power save features and sub-channelisation are all important elements necessary to gain momentum in the upcoming market for mobile broadband. Enhanced implementation of MIMO and adaptable antenna systems are also necessary in order to achieve higher throughput without sacrificing bandwidth. In addition, the new amendment enhances some of the shortcomings of the old standard, improves the security, adds support for frequency reuse, and adds another class for QoS.

The propagation modelling that was carried out in this project seems to represent the measured coverage quite well. Some deviation was found, where the measured signal strength was less than expected. The causes of this seems to be related to both differences in parameters used in the model and during testing, and lack of detail in the cartographic layers.

Testing along the tram line revealed that the combined coverage from the two planned sectors was only adequate to achieve high throughput along one third of the tram line. The base station that was located at site BS1 provided very good coverage up to about 1.5 kilometre in distance, and it is recommended that no changes is made this sector. A new sector pointing east-wards should also be added from this site to cover the remaining parts of the inner city.

The base station that was located at site BS2 did not provide as good coverage as expected. Measurements revealed that high throughput was only achieved up to 0.5 kilometres from the base station. This translates to less than one kilometre of the tram line. It is recommended to move this base station to a new location, in order to provide better coverage.

Testing of diversity showed that 2^{nd} order diversity did not improve the coverage area to a large degree. Improvements of a few hundred meters only were measured. Diversity using separation in space seemed to have a greater impact on throughput, where improvements of 42% and 52% were measured from site BS1 and site BS2 respectively. Separation by polarisation was only tested from site BS1, where improvements of 14% was measured.

One of the main problems that were found was related to speed. When the tram was in motion, average throughputs of more than 1 Mbps were seldom achieved. While standing still, average throughputs of 4 to 6 Mbps were easily measured. As the maximum throughput using a UL/DL ratio of 50/50 is 9 Mbps, this accounts to only 70% efficiency. The low throughput while in motion is suspected to be related to the use of an omni-directional antenna on the tram, which caused some degradation with regards to Doppler spread. It is recommended to make use of an array of sectorised antennas instead, in order to enhance the performance of the system.

8.2 Future Work

When mobile stations such as PCMCIA and USB implementations become available, testing should be carried out to investigate the indoor coverage provided.

WiMAX provides a number of QoS classes that may be used to provide different levels of services. How are they implemented and how do they perform with regards to voice and video streaming? How can they be integrated with QoS schemes already in use in Wireless Trondheim?

The antenna used on the tram was omni-directional. It seems that this antenna was not optimal with regards to a moving vehicle. It would be interesting to test other types of antennas with different gains and check for any improvements. It is especially recommended to test an array of sectorised antennas.

Perform Doppler shift measurements on the tram while in motion.

When devices with both WiMAX and Wi-Fi chipsets become available, testing should be carried out to verify how inter-media handover between the two access technologies work.

The WiMAX system may be managed using SNMP. Which management systems are available and how do they work? What data may be fetched from the base station? How can data be gathered and analysed in order to improve the quality of the network?

A new security framework was added to the IEEE 802.16e-2005 amendment. It would be interesting to gain more knowledge about the improvements and check for any vulnerabilities.

Implement 4^{th} order diversity and perform testing of coverage and throughput.

Perform new throughput measurements using the UDP protocol instead of TCP. It would be particularly interesting to investigate the throughput in the uplink channel with regards to the number of sub-channels in use, and if there is indeed any relationship between the number of uplink sub-channels in use and throughput of TCP in the downlink.

Investigate the bit error rates on the physical level and determine the causes of packet loss.

When the system has been upgraded to mobile WiMAX, handover between cells and handover between different operators should be investigated and tested.

Investigate how the physical WiMAX network may be used by several service providers, just like Wireless Trondheim offers different service providers access to a common physical network.

Perform more planning with the goal of covering the entire city of Trondheim with WiMAX access.

Bibliography

[1] Ball State University. WiMAX Technology Evaluation, 2007.

[2] O. Grøndalen, P. Grønsund, T. Breivik, and P. A. Engelstad. Fixed WiMAX Field Trial Measurements and Analyses. In Pål Grønsund, editor, *Mobile and Wireless Communications Summit, 2007. 16th IST*, pages 1–5, 2007.

[3] J. G. Andrews, A. Ghosh, and R. Muhamed. *Fundamentals of WiMAX: Understanding Broadband Wireless Networking*. Prentice Hall, 2007.

[4] IEEE. IEEE Standard for Local and Metropolitan Area Networks Part 16: Air Interface for Fixed Broadband Wireless Access Systems, 2002.

[5] IEEE. IEEE Standard for Local and Metropolitan Area Networks Part 16: Air Interface for Fixed Broadband Wireless Access Systems Amendment 2: Medium Access Control Modifications and Additional Physical Layer Specifications for 2-11 GHz, 2003.

[6] IEEE. IEEE Standard for Local and Metropolitan Area Networks Part 16: Air Interface for Fixed Broadband Wireless Access Systems Amendment 1: Detailed System Profiles for 10-66 GHz, 2002.

[7] IEEE. IEEE Standard for Local and Metropolitan Area Networks Part 16: Air Interface for Fixed Broadband Wireless Access Systems, 2004.

[8] IEEE. IEEE Standard for Local and Metropolitan Area Networks Part 16: Air Interface for Fixed and Mobile Broadband Wireless Access Systems Amendment 2: Physical and Medium Access Control Layers for Combined Fixed and Mobile Operation in Licensed Bands and Corrigendum 1, 2006.

[9] IEEE. IEEE Standard for Local and Metropolitan Area Networks Part 16: Air Interface for Fixed Broadband Wireless Access Systems Amendment 1: Management Information Base, 2005.

[10] IEEE. IEEE Standard for Local and Metropolitan Area Networks Part 16: Air Interface for Fixed and Mobile Broadband Wireless Access Systems Amendment 3: Management Plane Procedure and Services, 2007.

[11] IEEE. IEEE Standard for Local and Metropolitan Area Networks Media Access Control (MAC) Bridges Amendment 5: Bridging of IEEE 802.16, 2007.

[12] L. Nuaymi. *WiMAX: technology for broadband wireless access*. Wiley, Chichester, 2007.

[13] WiMAX Forum. The WiMAX Forum Certified Program for Fixed WiMAX (Whitepaper), 2007.

[14] WiMAX Forum. Mobile WiMAX Part I: A Technical Overview and Performance Evaluation (Whitepaper), 2006.

[15] ISO. Standard 7498-1: Information Technology - Open Systems Interconnection - Basic Reference Model: The Basic Model, 1994.

[16] IEEE. IEEE Standard for Local and Metropolitan Area Networks Part 2: Logical Link Control, 1998.

[17] WiMAX Forum. Mobile WiMAX System Profile Release 1.0 Revision 1.4.0, 2007.

[18] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(July and October):379–423 and 623–656, 1948.

[19] C. Berrou and M. Jezequel. Non-binary convolutional codes for turbo coding. *Electronics Letters*, 35(1):39–40, 1999. 0013-5194.

[20] F. Gray. Pulse Code Communication, March 17 1953.

[21] R. Olexa. *Implementing 802.11, 802.16, and 802.20 wireless networks: planning, troubleshooting and operations*. Elsevier, Amsterdam, 2005. 1 CD-ROM (12 cm) i lomme.

[22] R. W. Chang. Synthesis of band-limited orthogonal signals for multichannel data transmission. *Bell System Technical Journal*, 46:1775–1796, 1966.

[23] Cable Television Laboratories. Data-Over-Cable Service Interface Specifications (DOCSIS 1.1) - Baseline Privacy Plus Interface Specification, 2005.

[24] D. Johnston and J. Walker. Overview of IEEE 802.16 security. *IEEE Security & Privacy Magazine*, 2(3):40–48, 2004.

[25] P. Frenger, S. Parkvall, and E. Dahlman. Performance comparison of HARQ with Chase combining and incremental redundancy for HSDPA. In S. Parkvall, editor, *Vehicular Technology Conference, 2001. VTC 2001 Fall. IEEE VTS 54th*, volume 3, pages 1829–1833 vol.3, 2001.

[26] J. F. Cheng. Coding performance of hybrid ARQ schemes. *Communications, IEEE Transactions on*, 54(6):1017–1029, 2006. 0090-6778.

[27] Motorola. A Practical Guide to WiMAX Antennas: MIMO and Beamforming Technical Overview (Whitepaper), 2007.

[28] A. Goldsmith. *Wireless Communications*. Cambridge University Press, 2005.

[29] J. R. Barry, E. A. Lee, and D. G Messerschmitt. *Digital Communication*. Springer, 2003.

[30] B. Baumgärtner. Illustration of differentiation between SISO, SIMO, MISO and MIMO wireless communication systems, 2005.

[31] S. M. Alamouti. A simple transmit diversity technique for wireless communications. *Selected Areas in Communications, IEEE Journal on*, 16(8):1451–1458, 1998.

[32] WiMAX Forum. WiMAX Forum Network Architecture Stage 2 and 3 Release 1 Version 1.2, 2008.

[33] WiMAX Forum. Mobile WiMAX: A Performance and Comparative Summary, 2006.

[34] Alvarion. Comparing Mobile WiMAX, 3G and Beyond (Whitepaper), 2007.

[35] W. Bolton, X. Yang, and M. Guizani. IEEE 802.20: Mobile Broadband Wireless Access. *Wireless Communications, IEEE*, 14(1):84–95, 2007.

[36] B. M. Bakmaz, Z. S. Bojkovic, and D. A. Milovanovic. Mobile Broadband Networking Based on IEEE 802.20 Standard. In Z. S. Bojkovic, editor, *Telecommunications in Modern Satellite, Cable and Broadcasting Services, 2007. TELSIKS 2007. 8th International Conference on*, pages 243–246, 2007.

[37] S. H. Andresen, J. Krogstie, and T. Jelle. Lab and Research Activities at Wireless Trondheim. In *International Symposium on Wireless Communication Systems 2007*, Trondheim, 2007.

[38] Sesam. Sesam map, 2008.

[39] ATDI. ICS telecom nG, 2007.

[40] R. Coudé. Radio Mobile, 2008.

[41] U.S. Geological Survey. Shuttle Radar Topography Mission, 2007.

[42] NLANR. Iperf, 2008.

[43] P. Grønsund. A Field Study of WiMAX Performance, 2007.

[44] X. Ma, H. Kobayashi, and S. C. Schwartz. Effect of frequency offset on BER of OFDM and single carrier systems. In *Personal, Indoor and Mobile Radio Communications, 2003. PIMRC 2003. 14th IEEE Proceedings on*, volume 3, pages 2239–2243 vol.3, 2003.

[45] M. Russell and G. L. Stuber. Interchannel interference analysis of OFDM in a mobile environment. In *Vehicular Technology Conference, 1995 IEEE 45th*, volume 2, pages 820–824 vol.2, 1995.

[46] G. Zaggoulos, A. Nix, and A. Doufexi. WiMAX System Performance in Highly Mobile Scenarios with Directional Antennas. In *Personal, Indoor and Mobile Radio Communications, 2007. PIMRC 2007. IEEE 18th International Symposium on*, pages 1–5, 2007.

[47] J. Joe, S. K. Hazra, S. H. Toh, W. M. Tan, and J. Shankar. 5.8 GHz Fixed WiMAX Performance in a Sea Port Environment. In *Vehicular Technology Conference, 2007. VTC-2007 Fall. 2007 IEEE 66th*, pages 879–883, 2007.

[48] John P. Snyder. Map Projections: A Working Manual, 1987.

Appendices

.1 Electronic Attachments

The electronic attachments to this thesis include the following:

Directory **Alvarion** contains public data sheets and product information about the WiMAX solution that was used for testing

Directory **Graakallbanen** contains all the raw data that was obtained during testing, grouped into six folders. Day 1 contains data from the preliminary testing from BS1. Day 2 and 3 contain data from testing on the tram from BS1. Day 4 contains data from the preliminary testing from BS2, both phase I and phase II. Day 5 and 6 contain data from testing on the tram from BS2. Each group contain raw physical performance data fetched from the CPE along with raw GPS data and raw throughput measurement data. Additionally, Microsoft Excel spreadsheets contain all the data that was analysed and also the graphs that were made. Garmin Mapsource files contain GPS waypoints and tracks for all the tests.

Directory **Pictures** contains a number of photos grouped into five folders. The photos show the process of installing the base station at each site, show how the antennas were mounted, and where they were placed. A photo is also taken at each of the locations visited during the preliminary testing phases. The photo taken shows the view towards the base station.

Directories **PT February** and **PT April** contain all the modelling results from two visits to the Norwegian Post- and Telecommunications Authority. In addition to the results, screenshots of the different settings used are also included.

Directory **Source Files** includes source files of the software programs used during testing.

.2 Frequency Auctions

Even though IEEE 802.16 is specified for a wide frequency range, the WiMAX Forum only certifies products in a few defined bands. The Norwegian Post and Telecommunications Authority has conducted auctions for the frequencies relevant for WiMAX deployment in Norway. These are 2.3, 2.6, and 3.5 GHz. The results of the auctions are listed below in chronological order.

Region	Population	Percentage
1	1.957.946	42%
2	673.027	14%
3	808.290	17%
4	371.729	8%
5	407.905	9%
6	462.237	10%
Total	4.681.134	100%

Table 1: Population per region

.2.1 3.5 GHz

The auction for the 3.5 GHz frequency band was held in December 2004. The entire country was divided into 6 regions, as illustrated in figure 1. There were allocated 25 slots per region, where each slot except the last one is a 2x3.5 MHz paired channel suitable for FDD. In total, 173 MHz was to be auctioned. The licence expire at the end of 2022, which gives them a lifetime of 18 years.

The auction lasted for 8 days and received a total of approximately 50 MNOK. 8 companies secured frequency slots, where UPC, Telenor, and NextGenTel together won 75% of the available frequency spectrum. Catch, Netpower, Hardanger, Nera, and BaneTele won the remaining 25%. The average price paid per slot in each region is outlined in figure 2. As can be seen, region 3 achieved the highest prices with an average of close to 750.000 NOK per slot. The overall average price paid per slot was 330.000 NOK.

It is also interesting to investigate the price paid per MHz in relation to the population. These figures are calculated as the average price paid in NOK divided by the total available MHz and divided again by the population in the desired area. The population used in the calculation is listed in table 1. A graphical illustration of these figures are presented in figure 3. The average cost is 0.062 NOK, while the cost for region 3 is more than twice the average.

Figure 4 illustrates the price paid per slot for each company and region. As seen from the figure, Telenor, Catch, and NextGenTel secured frequency spectrum in all regions. On the other hand, some companies only bid in selected regions.

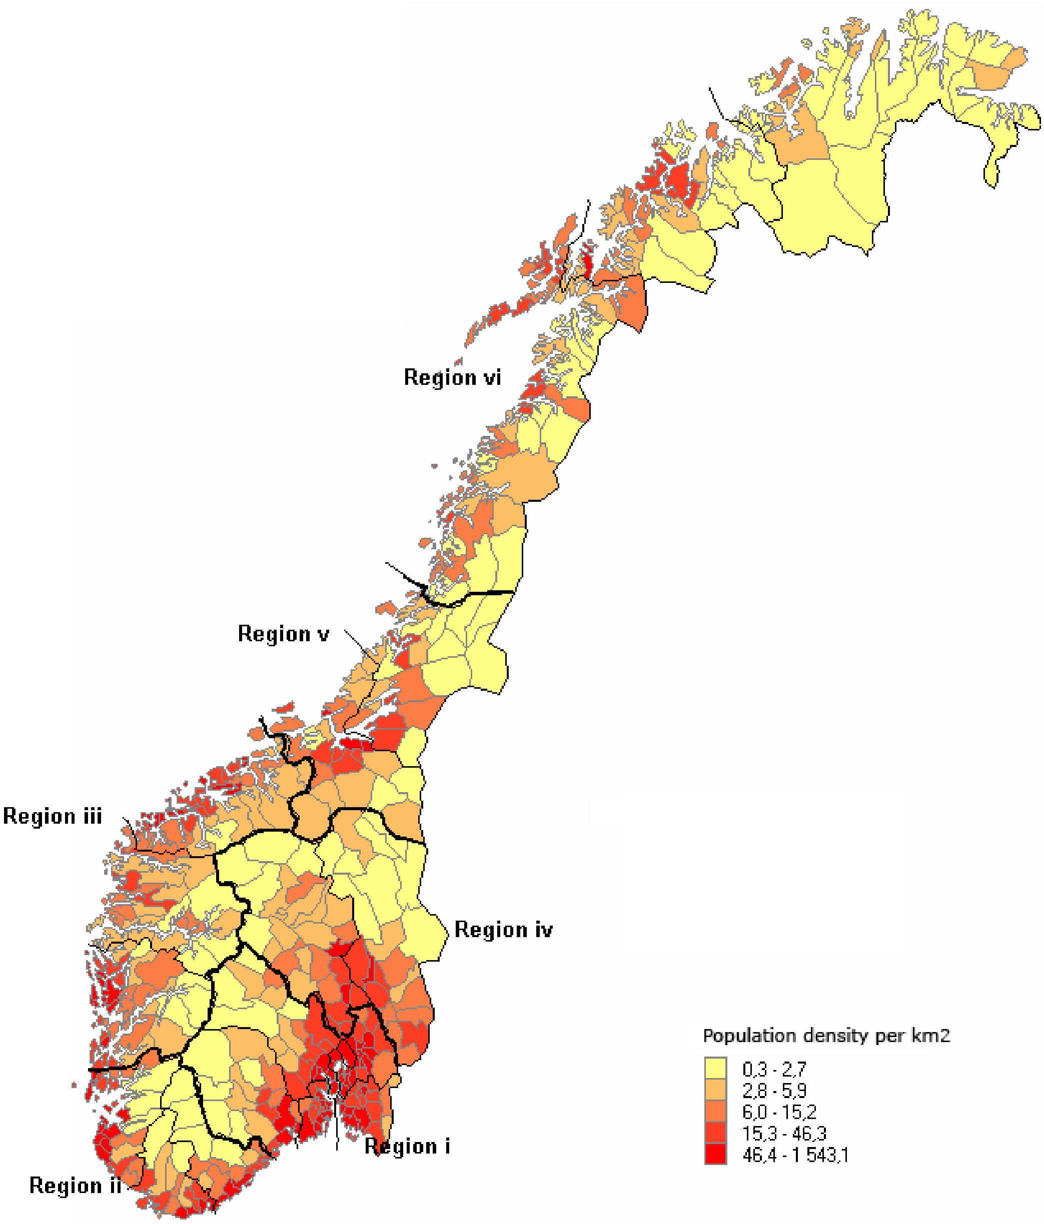

Figure 1: Map of regions

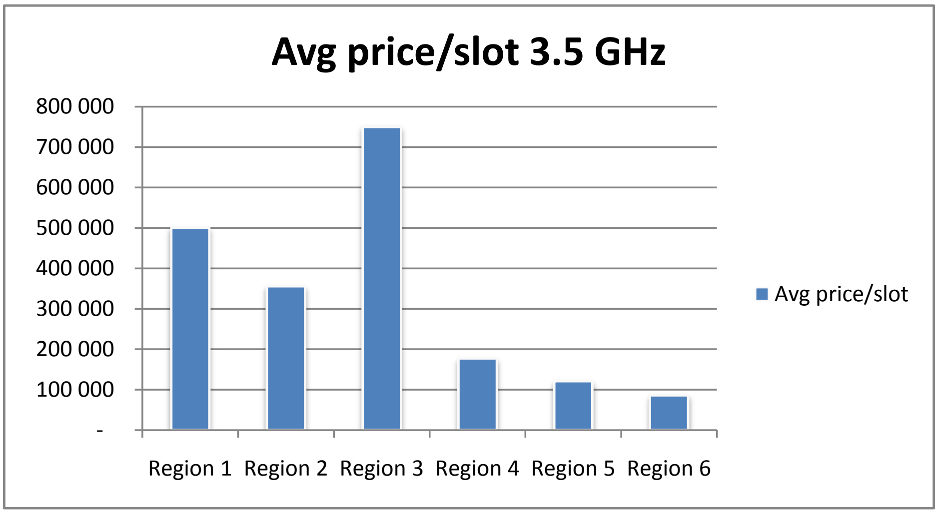

Figure 2: Average price/slot per region for 3.5 GHz

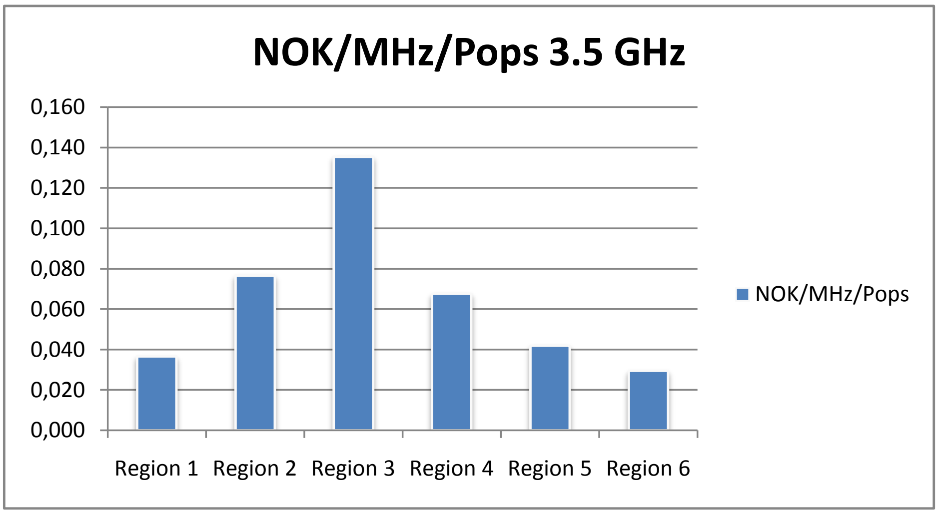

Figure 3: NOK/MHz/Pops for 3.5 GHz

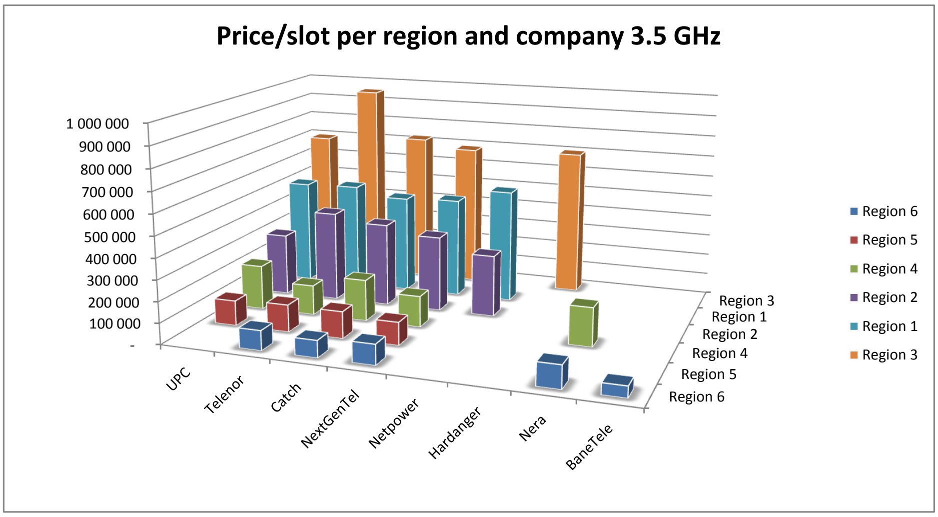

Figure 4: Price/slot per region and company for 3.5 GHz

.2.2 2.3 GHz

The auction for the 2.3 GHz band was held in September 2006. It consisted of a single 22 MHz slot covering the whole country, hence it was not divided into regions as with the previous auction. A total of 8 companies participated in the bidding, which lasted one day. At the end, NextGenTel won the auction with a total price of 7 MNOK.

The average NOK/MHz/Pops for this slot was 0.068 NOK, which is in the same range as the auction for 3.5 GHz.

.2.3 2.6 GHz

The auction for the 2.6 GHz band was held in November 2007. A total of 190 MHz was bid for over the duration of seven days, and the total income amounted to approximately 131 MNOK. Eight companies participated in the bidding process, while only five actually secured spectrum.

Most of the spectrum was bought by Telenor, Netcom, and Craig Wireless. They all received spectrum in all of the six regions. The two remaining companies, Hafslund and Arctic Wireless, received spectrum in some of the regions. At the end of the auction, three slots were not sold. They were auctioned once more in February 2008. Hafslund delivered the highest bid and received more spectrum licences in regions 2, 3, and 5.

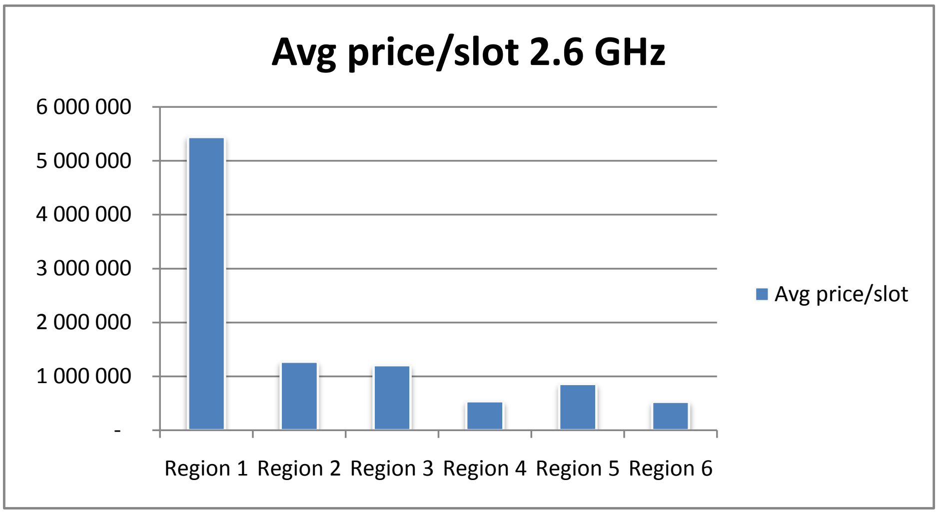

Figure 5: Average price/slot per region for 2.6 GHz

Figure 5 outlines the average price per slot for the auction. It is clear that region 1 obtained the highest prices, with an average of close to 5.5 MNOK per slot. The other regions received more or less similar prices, varying between 1.2 MNOK and 0.5 MNOK.

The NOK/MHz/Pops ratios are illustrated in figure 6. Region 1 achieved the highest ratio with close to 0.353 NOK. This is close to a six-time increase compared to the two former auctions. On average, the NOK/MHz/Pops ratio is 0.259 for all regions. This number is also considerably higher than what was received earlier. One reason for this is that the 2.6 GHz band is very attractive with regards to mobile communication, as it is available for both WiMAX and other 3G technologies.

Figure 7 illustrates the price paid per slot for each company and region. As seen from the figure, Telenor, NetCom, and Craig Wireless secured frequency spectrum in all regions.

.2.4 Comparison

A comparison of the NOK/MHz/Pops ratios achieved for all three auctions is laid out in figure 8. The 2.3 GHz auction consisted of a single block, hence its ratio is equal for all regions. It is interesting to note that the 2.6 GHz auction received the highest ratio in region 1, while the 3.5 GHz auction received the highest price in region 3.

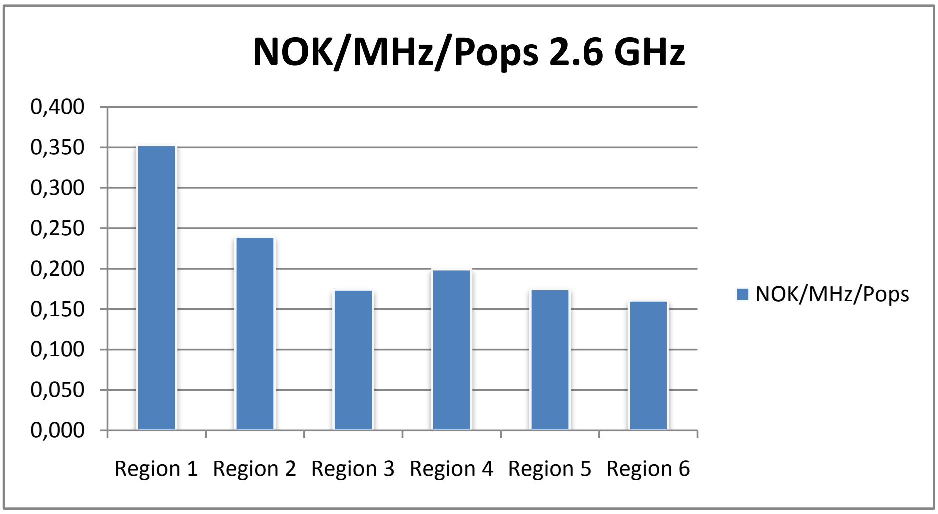

Figure 6: NOK/MHz/Pops for 2.6 GHz

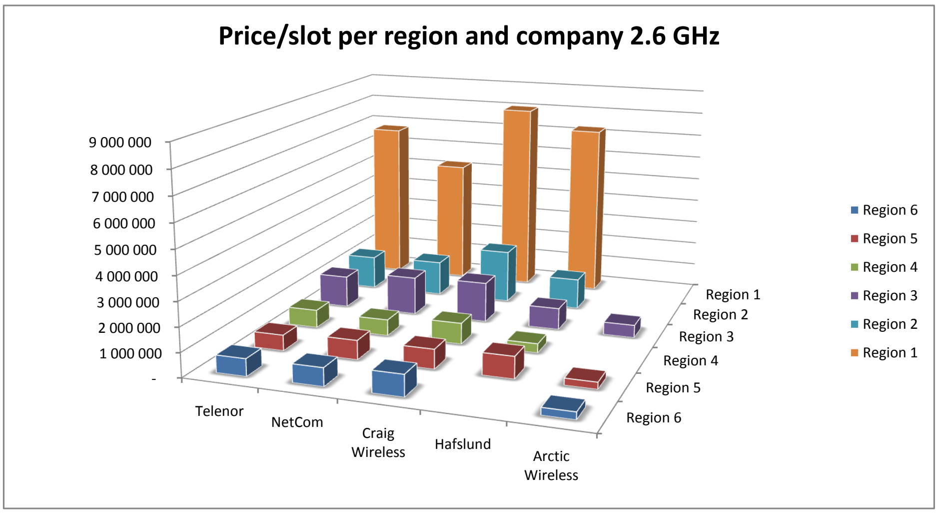

Figure 7: Price/slot per region and company for 2.6 GHz

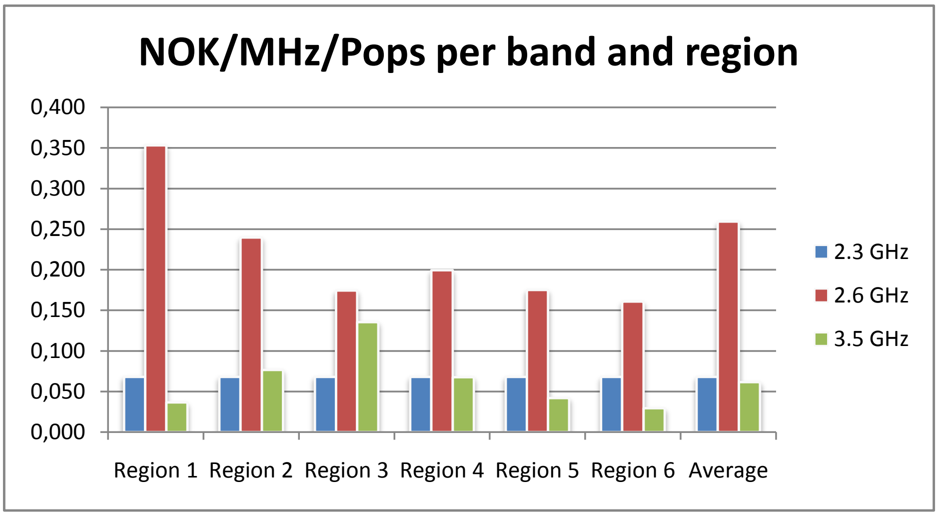

Figure 8: NOK/MHz/Pops comparison for the three auctions

.3 Script used for Monitoring the CPE

The following Python script was written in order to continuously monitor and write the physical performance parameters from the CPE to a file.

```
#!C:\Program Files\Python25

import os
import telnetlib
import sys
import datetime
import string
import time
from time import localtime, strftime

TIME = strftime("%H%M%S", localtime())
IP_CPE = "192.168.254.251"
PASSWORD = "installer"
ESC = chr(27)
ITERATE = 1200

# Open file for writing
outfile = open("LinkQuality_" + TIME +".txt","wb")
```

```
# Create a new telnet connection to the CPE
telnet = telnetlib.Telnet(IP_CPE)

# Authenticate the connection
telnet.read_until("Enter the password: ")
telnet.write(PASSWORD + "\r\n")
telnet.read_until("Enter your choice > ")
telnet.write("7\r\n")
telnet.read_until("Enter your choice > ")

print "Starting the link quality display...\r\n"

# Start the link quality display
telnet.write("1\r\n")
telnet.read_until("Started. Press ESC to stop...")

# Loop while fetching one line at the time
for i in range(1, ITERATE + 1):
LINK_DATA = telnet.read_until("\n")
CUR_TIME = strftime("%H:%M:%S", localtime())
outfile.write(CUR_TIME + " " + LINK_DATA)
print CUR_TIME + " " + LINK_DATA ,
else:
# Clean up and exit
telnet.write(ESC)
telnet.write("9\r\n")

telnet.close()
outfile.close()
```

.4 Distance Calculation

At sea level, the length of a latitudinal minute is constant and measures approximately 1852 meters all over the globe. The width of a longitudinal minute is not constant and depends on the latitude, hence it must be calculated. The reason for this is that the degrees of longitude converge towards the poles, making the distance between them less towards the poles than at the equator. The width of a latitudinal degree is calculated using the following formula [48]:

$$\frac{\pi}{180^\circ}\cos(\phi)\sqrt{\frac{a^4\cos(\phi)^2+b^4\sin(\phi)^2}{(a\cos(\phi))^2+(b\sin(\phi))^2}} \tag{1}$$

where ϕ is the degree of latitude, and a and b are the Earth's equatorial and polar radii equal to 6,378,137 m and 6,356,752.3142 m respectively when using the WGS84 datum. As the tram line lies at approximately 63 degrees 25 minutes north, ϕ equals $63 + 25/60 \approx 63.42$. Using this value yields the length of a latitudinal minute at this parallel equal to 828 m.

In order to calculate the Euclidean distance between any two points in a Cartesian coordinate system, the Pythagoras' theorem may be used. Note that this method calculates the distance of a straight line between the two points, which is not entirely correct as the curve of the Earth is not taken into account. This error is however negligible.

$$d = \sqrt{((x_1 - x_2) * 828)^2 + ((y_1 - y_2) * 1852)^2} \tag{2}$$

Using the above formulae, the distance between any two neighbouring coordinates was calculated and added to the total track length.