

On Reliability in Mobile Broadband Networks

Džiugas Baltrūnas

July 2017

Thesis submitted for the degree of Philosophiæ Doctor

Liudvikui ir Gertrūdai

Abstract

Mobile broadband (MBB) networks are arguably becoming the most important component in the modern communications infrastructure. Their widespread popularity has given them a role as *critical infrastructure*. Failures in MBB networks have large consequences for users, and large-scale outages receive significant attention in the media and from governing bodies. The importance of MBB networks demands quantification of their stability and performance, as experienced by the end user.

Experienced quality is a complex notion, which encompasses different performance metrics. The reliability must be measured at several levels, from the stability of the network connection to the reliability of the data plane and application layer performance. In this thesis, we present a framework to measure the reliability at control, data and application levels. We then measure and quantify the experienced reliability in Norwegian MBB networks, and compare reliability between networks. We assess the quality and reliability experienced by end users through systematic end-to-end measurements. The measurements are performed using Nornet Edge (NNE) – a large-scale dedicated measurement infrastructure that consists of more than hundred measurement nodes each connected to two or more MBB networks.

The studies presented in this thesis look at MBB reliability and performance in two different scenarios: when MBB connections are stationary, and when they are moving. We observe clear differences in the reliability of stationary and moving connections, and between operators. Mobile connections experience severe packet loss and are also exposed to a set of challenges not common for the stationary scenario. We find that MBB networks vary in the stability of connections, in packet loss and delay patterns, and in their ability to support popular applications. We also show that using two MBB connections from distinct operators in parallel can potentially give 99.999% availability.

The last part of the thesis presents MULTEX – a novel approach that enables

mobile devices to have simultaneous access to multiple packet data networks and thus increase the reliability by built-in multi-homing feature of Long Term Evolution (LTE).

Acknowledgements

The process of understanding reliability of mobile broadband networks throughout my PhD was not always reliable. There were interrupts, delays and outages resulting in practical and psychological challenges of different scale. Addressing and solving those challenges and, most importantly, learning from them would not have been possible without help and support from others.

Amund Kvalbein, my former principal supervisor, believed in my ability to pursue a PhD from the very beginning. He had so much trust in advance in me that once he dared to take me to a meeting with clerks in Norwegian a few months after I started and knew only *ha det bra*. His exceptional leadership skills and ability to inspire others to work on (sometimes crazy) ideas let anyone around him feel they were on the right track.

Ahmed Elmokashfi, my principal supervisor, took over Amund when he left Simula in 2014. Ahmed not only ensured the successful continuation of the project and my PhD, he was the best college I ever had. He had enormous personal capacity to listen to others patiently and advice on any situation or matter.

Professor Olav Lysne, my supervisor, did an amazing job in promoting the project nationally and shielding it from various threats when that was necessary. I do not remember any meeting with Olav without him mentioning how important it was what we did. His encouragement and versatile support was invaluable to me personally and for the overall success of the project.

I thank Simula for providing a splendid multicultural environment and an opportunity to meet and work with so many great colleagues. I also thank all the partners who supported the project and our collaborators from all over the world.

Last but not least, I do not think I can ever repay a debt to my wife Živilė. She quit her job where she had excellent career opportunities and left with me to Norway because of me. She gave birth to our two wonderful children and raised them while I was doing my PhD. She never stopped believing in me and supporting me. *Ačiū!*

Contents

1	Introduction	7
1.1	The increasing importance of mobile broadband networks	7
1.2	Demand for information on MBB performance and stability	8
1.3	The need for end-to-end measurements using a dedicated infra- structure	9
1.4	Contributions	10
1.5	Publications	11
I	A framework and a platform for measuring mobile broad- band reliability	13
2	The evolution and architecture of MBB networks	15
2.1	The basic functions and the evolution of mobile networks	15
2.1.1	2G	15
2.1.2	3G	21
2.1.3	4G	22
2.1.4	5G	23
2.2	MBB network architecture	24
2.2.1	RAN	26
2.2.2	RRC state machine	26
2.2.3	PDP context, EPS bearer and data tunnelling	28
2.2.4	CN components	28
3	Measuring MBB network reliability	31
3.1	The definition and importance of reliability	31
3.2	Reliability metrics	32
3.3	Proposed framework for measuring reliability	32

3.3.1	Connection level reliability	33
3.3.2	Data plane reliability	34
3.3.3	Application level reliability	35
3.4	Related approaches for measuring MBB networks	35
3.4.1	Platforms with dedicated measurement nodes	37
3.4.2	Crowd-sourcing platforms and applications	39
3.4.3	Measurements using network-side data	43
3.4.4	Other approaches to measure MBB networks	45
3.4.5	Related MBB reliability and performance studies	45
3.5	Summary	46
4	The Nornet Edge platform for mobile broadband measurements	49
4.1	Essential features	49
4.2	System overview	50
4.2.1	Deployment of nodes	51
4.2.2	Multi-homing	52
4.3	The Nornet Edge node	52
4.3.1	Technical specifications	54
4.3.2	Management and support functions	56
4.4	Backend system	59
4.4.1	Node control and management	59
4.4.2	Database	61
4.4.3	Visualisation	62
4.5	Nodes on trains	63
4.6	Challenges with running a large-scale measurement infrastructure	64
4.6.1	Maintaining high number of active nodes and MBB sub- scriptions	64
4.6.2	Typical failures and preventive actions	65
4.7	Summary	67
II	A large-scale study of reliability of Norwegian MBB net- works	69
5	Measured MBB networks and measurement datasets	71
5.1	Measured MBB networks	71
5.2	The effect of mobility on MBB reliability and performance	72
5.3	Measurement datasets	73

5.3.1	Stationary scenario	74
5.3.2	Mobile scenario	74
6	Reliability and performance: stationary scenario	77
6.1	Connection reliability	77
6.1.1	Measuring connection failures	78
6.1.2	Analysing connection failures	80
6.1.3	Correlating with metadata	83
6.1.4	Summary of findings	86
6.2	Data plane: measuring packet loss	86
6.2.1	Measurement description and scenarios	86
6.2.2	Loss rate and loss runs	87
6.2.3	Capturing large failure events	90
6.3	Data plane: dissecting packet loss	92
6.3.1	Packet loss statistics	93
6.3.2	Correlations in loss	99
6.3.3	Operator-wide correlations	99
6.3.4	RNC-level correlations	101
6.3.5	Cell-level correlations	103
6.3.6	Possible causes	104
6.3.7	Discussion	108
6.3.8	Summary of findings	110
6.4	Data plane: investigating excessive delays	110
6.4.1	Basic delay statistics	111
6.4.2	Characterising excessive delays	114
6.4.3	Underlying causes	120
6.4.4	Summary	121
6.5	Application level reliability	122
6.5.1	HTTP download test	122
6.5.2	VoIP test	124
6.5.3	Summary of findings	125
6.6	Robustness by multi-homing	125
6.6.1	Correlations in coverage	126
6.6.2	Correlations in downtime	127
6.6.3	Correlations in loss	128
6.6.4	Summary of findings.	129
6.7	Summary	130

7	Reliability and performance: mobile scenario	131
7.1	Classification of loss under mobility	131
7.1.1	Measurement setup	132
7.1.2	Measurement scenario	134
7.1.3	Data curation	135
7.1.4	Loss in different RATs	136
7.1.5	Classifying loss under mobility	137
7.1.6	Varying RAT	139
7.1.7	Constant RAT	147
7.2	Investigating excessive delays under mobility	150
7.2.1	Basic delay statistics	151
7.2.2	Frequency of triangle events	151
7.2.3	Summary of findings	153
7.3	Summary	153
8	Evolution of MBB networks	157
8.1	Shorter and less frequent downtimes	157
8.2	Significant reduction in loss	160
8.3	Fewer and less severe large failure events	162
8.4	Summary	164
III	Towards more reliable and flexible networks	165
9	Routing limitations in current LTE networks	167
9.1	The problem of inflexible routing	167
9.2	The support for multiple PDN connectivity in current LTE standards	169
9.2.1	EPS bearers and tunnels	169
9.2.2	Traffic offloading	169
9.2.3	Multiple PDNs	169
9.3	Shortcomings with the mobile network architecture	170
9.3.1	Suboptimal routing due to centralisation	170
9.3.2	Statefulness	170
9.3.3	Lack of geographically limited networks	171
9.3.4	Lack of reliability through built-in diversity	171
9.4	Related work	171
9.4.1	LIPA and SIPTO	172
9.4.2	NFV and SDN based initiatives	172

9.4.3	Other initiatives	173
10	MULTEX: towards more flexible routing and services in mobile networks	175
10.1	Motivation	175
10.2	Design and philosophy	176
10.3	Required changes	177
10.4	Use cases	178
10.5	Road towards implementation	179
10.5.1	Discovery of available PDNs	179
10.5.2	Routing and decision models	182
10.6	A proof of concept of MULTEX	183
10.6.1	Scenario	183
10.6.2	Local gateways and PDNs	184
10.6.3	LTE attach procedure	185
10.6.4	PDN announcement inside the PCO	185
10.6.5	Additional PDNs parameters	185
10.6.6	Additional PDN discovery and handling by the UE	186
10.7	Evaluation	186
10.7.1	Support for multiple default EPS bearers in current LTE networks and UEs	187
10.7.2	Performance of multiple default EPS bearers in real LTE networks	188
10.7.3	The cost of maintaining simultaneous EPS bearers	188
10.7.4	Charging, security and mobility.	190
10.8	Discussion and conclusions	191
IV	Epilogue	193
11	Conclusions	195
	Acronyms	197
	Bibliography	207

List of Figures

2.1	The evolution of mobile networks.	16
2.2	GPRS protocol stack[29, 40].	19
2.3	Simplified GSM/UMTS MBB network architecture.	25
2.4	Simplified LTE network architecture.	25
2.5	RRC state machines of GSM, UMTS and LTE networks.	27
3.1	Framework for measuring experienced reliability in MBB networks.	33
4.1	NNE overview.	50
4.2	Placement of NNE nodes in Norway as of June 2014.	53
4.3	NNE node with four USB modems.	55
4.4	The most important functions on the NNE node.	56
4.5	NNE backend.	59
4.6	Screenshot from the NNE visualisation website.	63
4.7	The number of active NNE nodes and MBB connections over time. Note that the counters for Tele2 and Ice connections drop to zero in 2015 because of the shut down of the respective networks.	64
5.1	The operators and RANs measured in this study.	71
5.2	Overall loss rate and the effect of mobility. Much higher loss rates observed when mobile in both networks.	73
5.3	Measurement datasets used in the study.	74
6.1	The statistics of connection failures.	79
6.2	The CDF of the average RSSI per operator.	82
6.3	The daily median downtime percentage for each MBB operator.	83
6.4	Downtime correlation with connection RAT, RSSI and E_c/I_o	84
6.5	Loss rate for each MBB operator.	87
6.6	Median daily loss rate for each MBB operator.	88
6.7	Loss, downtime and signal quality.	89

6.8	The distribution of loss run sizes across operators.	90
6.9	Large failure events July 2013 – May 2014.	91
6.10	Loss rate for each MBB operator. We see clear differences between operators.	93
6.11	Loss rate experienced during different RRC bins. We record very high loss rate during periods with state transitions.	97
6.12	Loss rate for each MBB operator considering only non-pathological states.	98
6.13	Correlations for loss $\geq 3\%$ in different RRC categories for Telenor’s RAN (top) and Netcom’s RAN (bottom). Loss in DCH and FACH exhibits relatively strong correlations.	100
6.14	Correlation in loss for connections served by the same RNC vs connections served by different RNCs.	102
6.15	Correlation in loss for connections served by the same cell vs connections served by different cells.	104
6.16	Loss rate quantification for Telenor and Netcom MBB networks. .	105
6.17	Diurnal patterns of loss in Telenor’s RAN (top) and Netcom’s RAN (bottom). DCH+FACH loss climbs rapidly during peak hours. . .	107
6.18	Loss rate by RRC state (top) and median loss rate (bottom) for Telenor after the RNC reconfiguration.	109
6.19	Average and max RTTs for stationary connections on Telenor (one-minute bins). Except for bins with mixed-RATs and 2G, RTTs are stable. 3G-DCH and LTE maximum RTTs are double the averages. .	112
6.20	Distribution of maximum RTT coefficient of variation for Telenor connections. LTE, 3G-DCH, and 2G connections have highly variable maximum RTTs.	113
6.21	Triangle event illustration.	115
6.22	The distribution of the frequency of triangle events for Telenor (top) and Netcom (bottom). The solid line marks the fraction of connections that experience more than a single event per hour on average. Approximately one quarter of 3G-DCH connections experience one or more events per hour, on average.	117
6.23	PDF of triangle event duration for Telenor (top) and Netcom (bottom).	118
6.24	Median packet inter-arrival at server during triangle events for Telenor (left) and Netcom (right).	119
6.25	Failure rates in HTTP download tests.	123

6.26	Probability of achieving less than 1Mbps HTTP throughput. . . .	124
6.27	Failure rates in VoIP tests.	125
6.28	Average E_c/I_o values for connections in different operators.	126
6.29	Common downtime across operators.	127
6.30	Conditional probability of loss events.	129
7.1	Map of the train routes.	133
7.2	Typical sequence of connectivity and coverage conditions that MBB connections experience as they move.	134
7.3	Loss rate for 3G, LTE and when a RAT change is involved (mixed). Most loss happens in mixed category, LTE performs the best. . . .	137
7.4	Classification of loss. The numbers given are percentages of lost packets relative to the parent category. In total, there are 63837 5-minute bins in the dataset. About 229992 packets were lost during these bins.	138
7.5	The percentages of lossy and non-lossy 5-minute bins for constant and varying RATs.	138
7.6	The percentages of lossy and non-lossy 5-minute bins with varying RATs split by whether more than one RAT is observed.	140
7.7	The percentages of lossy and non-lossy 5-minute bins for varying RAT with a RAT change split by the presence of connection resets. . . .	141
7.8	Involved RATs for the 5-minute bins with a RAT change and no connection resets. Clear differences between networks. More loss during 2G/3G handovers.	142
7.9	Involved Radio Access Technologies (RATs) for the 5-minute bins with a RAT change and one or more connection reset. Much more loss compared scenario without disconnects, minimal differences between networks, least lossy are bins with 3G/LTE handovers. . . .	143
7.10	Loss runs for involved RATs when there is a RAT change and one or more connection reset. Several modes around 10 packets in both networks. Different modes for different sets of involved RATs.	144
7.11	The percentages of lossy and non-lossy 5-minute bins for varying RATs with no RAT change split by the presence of connection resets. . . .	144
7.12	Individual varying RATs for the 5-minute bins when RAT does not change and there are no connection resets. Less loss in 3G bins with varying RAT in Telenor compared to Netcom.	145

7.13	Individual varying RATs for the 5-minute bins when RAT does not change and there is one or more connection reset. Differences between networks are minimal, but there is slightly less loss in 3G bins with varying RAT in Telenor compared to Netcom.	145
7.14	Loss runs for individual varying RATs when there is no RAT change, but one or more connection reset in a 5-minute bin intersecting with the loss run. All modes in a shorter range from 11 to 16, with a peak at 13 packets for both networks.	146
7.15	The percentages of lossy and non-lossy 5-minute bins for constant RAT split by the LAC change.	147
7.16	The percentages of lossy and non-lossy 5-minute bins for constant RATs split by the LAC and cell changes.	148
7.17	Loss rate in a 5-minute bin for 3G and LTE RATs split by the presence of one or more LAC change or cell change, or none of them. In both networks the highest loss occurs when there is a Location Area Code (LAC) change involved. In Telenor, bins with LAC changes are more lossy compared to Netcom, whereas in Netcom there is much loss during cell changes. When there is no LAC or cell change or the LTE cell changes, the loss is minimal in both networks.	149
7.18	Average and max RTT for mobile connections on Telenor (one-minute bins). 40% of 3G-DCH and 3G-FACH bins exhibit maximum delays ≥ 1 second.	152
8.1	The fraction of connections with MTBF < 1 day.	158
8.2	The fraction of connections with MTTR > 5 minutes.	158
8.3	The fraction of connections with the average downtime > 10 minute per day.	159
8.4	Median loss rate evolution in Telenor (top) and Telia (Netcom) (bottom) over time.	161
8.5	Median loss rate evolution.	162
8.6	Large failure events from July 2013 until December 2015.	163
10.1	Conceptual architecture of MULTEX.	176
10.2	DNS based PDN discovery.	181
10.3	A proof of concept of MULTEX.	184
10.4	Three parallel TCP downloads over a single (a) and three parallel EPS bearers (b).	189

List of Tables

6.1	Time spent (TS) and loss induced (LI) by each network on different RRC categories.	95
7.1	Fraction of lossy bins for Netcom cell and LAC changes for different speed (S) categories. The speeds are in km/h.	150
7.2	Triangle event frequency for mobile nodes.	153
10.1	Support for multiple default EPS bearers by different UEs.	187

Chapter 1

Introduction

1.1 The increasing importance of mobile broadband networks

MBB networks are arguably becoming the most important component in the modern communications infrastructure. The immense popularity of mobile devices like smartphones and tablets, combined with the availability of high-capacity third generation (3G) and fourth generation (4G) mobile networks, has radically changed the way most people access and use the Internet. The use of MBB networks has exploded over the last few years. Mobile data traffic has grown 4,000-fold over the past 10 years, and in 2015 alone, global mobile data traffic grew 74 percent [24]. An increasing number of people rely on their MBB connection as their *only* network connection, replacing both a fixed broadband connection and the traditional telephone line.

With increased use comes increased dependence. The popularity of MBB networks has given them a role as *critical infrastructure*. Organisations, businesses and individuals rely on a stable and well-provisioned MBB network in their day-to-day operations. MBB networks are used for a wide range of diverse tasks such as fleet management, connecting mobile payment terminals, real-time public transport information, safety alarms, and emergency warning systems. This dependence, primarily triggered by the growing numbers of smartphones, Internet of Things (IoT) devices and machine-to-machine (M2M) communication, is going to increase even further in the next generation MBB networks. Failures in MBB networks have large consequences for users, and large-scale outages receive significant attention in the media and from governing bodies. Examples of large events

include the fire in Rotterdam in April 2012 that affected millions of Vodafone customers [3], the software bug that took down Telenor’s network in Norway in June 2011 [2], the short-circuit in the power system of Digi’s network in Malaysia in November 2015 [64], and the national overnight outage of T-Mobile’s high-speed LTE data wireless in USA in September 2016 [34]. There are, however, a large number of smaller events that affect the user experience every day, which never reach the headlines.

1.2 Demand for information on MBB performance and stability

Given the importance of MBB networks, there is a strong need to quantify their stability and performance. In particular, it is important to measure and understand the *quality as experienced by the end user*. Experienced quality is a complex notion, which manifests itself in many different performance metrics. Can a data connection be established when it is needed? Can a data connection be maintained without disruption for a reasonably long time? What network capacity a user can expect? And what latency? Is it good enough to ensure reasonable Quality of Experience (QoE) of a Voice-over-IP (VoIP) session? Is it possible to stream a football match without interruptions?

Such performance information is sought by many parties. Regulators need it to control how operators fulfil their obligations, and as a basis for regulatory policies. Operators can augment their monitoring data with user-side measurements, in order to identify problems in their networks and reduce customer complaints. Organisations and businesses can use such information to build robust services and products on top of MBB networks. Providers of mobile services that run on top of MBB networks need reliable data on reliability in order to predict the performance of their own services. Application developers need to know the characteristics of the underlying network in order to build efficient applications. And finally, consumers can use such information to make informed choices on which network provider to choose.

There is, however, very little publicly available information today about the performance of MBB networks. For consumers, the only source of information is typically a network operator’s website that usually contains a coverage maps and historical reports about outages. In some cases, the coverage maps are based on data collected from drive tests, which some operators periodically conduct. In

other cases, they are based on theoretical calculations. Even regulators are often left with a posteriori event reports from the operators after large failure events without underlying technical details. The lack of reliable information about the MBB network performance led to several efforts and initiatives, for example, publicly accessible coverage and throughput measurements crowd-sourced from millions of smartphones.

1.3 The need for end-to-end measurements using a dedicated infrastructure

The ambition of this work is to *measure and quantify the experienced reliability in MBB networks*, and to compare reliability between networks. The quality experienced by end users can only be assessed through systematic end-to-end measurements. Monitoring from within the network can reveal parts of the picture, but can hardly capture the complex interplay between user equipment, network, protocols and applications. A working base station is not the same as a stable Internet connection, good signal strength is not the same as high throughput, light system load is not the same as stable delay and so on. Network-side data gives more insights into the network internals, while end-to-end measurements can help operators detecting anomalies that are not easily identifiable by only using network-side monitoring.

The approach presented in this thesis is based on building a dedicated infrastructure for measuring and experimenting in MBB networks. A dedicated measurement infrastructure can complement other approaches to measure reliability by overcoming the challenges discussed in Section 3.4. It can be used to run measurement experiments at regular intervals over long time periods under similar conditions. It also gives full information about the context in which the measurements are collected, and allows targeted experiments triggered by the observed behaviour.

As part of this thesis we present a framework for assessing MBB reliability using a dedicated measurement infrastructure with large scale deployment. We believe that this is a timely input to the ongoing efforts to defining approaches and metrics for measuring MBB networks. Our work is novel in that it uses a large, geographically distributed infrastructure, dedicated to recording MBB measurements.

Our study runs end-to-end measurements from a dedicated infrastructure consisting of several hundreds of measurement nodes. This not only gives us full control over all components of the measurement system, but also produces a dataset of better quality with fewer artefacts. We collect data from a dedicated infrastructure in order to have full control over the measurement nodes, allowing us to systematically measure the reliability over a long period of time.

1.4 Contributions

In summary, this thesis makes the following contributions:

1. We propose a *framework for measuring robustness in MBB networks*. The framework captures aspects of reliability on several layers, from a basic registration in the network to a stable application performance over time. Within this framework, we define metrics and measurement experiments that describe reliability on the connection level, the data plane level, and the application level.
2. We present the *first large-scale measurement study of MBB reliability*, from a dedicated measurement infrastructure. The measurement experiments are performed on NNE (Chapter 4). NNE is the largest infrastructure of its kind, with dedicated measurement nodes distributed in over 100 Norwegian municipalities and a set of mobile nodes. The data used in this work is captured from a over 900 MBB connections corresponding to more than 300 distinct nodes and 5 different operators. Through long-term monitoring of a large number of connections, we find that a significant but decreasing fraction of connections lose their network attachment more than 10 minutes per day. We also observe clear differences in reliability characteristics between networks: some connections experience failures more frequently but the their duration is shorter when compared to the infrequent but longer failures resulting in a higher overall downtime.
3. By capturing a rich set of metadata that describes the context of the measurements, this study *increases the value of end-user measurement data*. The metadata allows us to explain measurement results by looking at factors such as signal quality, radio state, network attachment, connection RAT, etc. In many cases, we are also able to distinguish between problems in the radio access network and the mobile core network. We find a clear correlation

between signal conditions, connection failures and loss, but we also discover that many packet loss or connection termination episodes can not be explained by signal quality alone. We further find that the inability to obtain dedicated radio resources is a common cause of application failures in some networks.

4. Measurements conducted from nodes placed on regional and inter-city trains captures MBB reliability and performance under mobility as experienced by the user. Measurement results from both stationary and mobile nodes lets us effectively compare quantify the impact of mobility and coverage on the performance under the two scenarios.
5. Geographical diversity of the measurement nodes gives a better understanding about the affected parts of the MBB network. Depending on the density of NNE nodes in a certain region, we are able to tell whether there is a correlation in failures at the cell, Radio Network Controller (RNC) or Core Network (CN) level. In one case, such correlations let us identify problems with the configuration that was shared across all RNCs of a single operator.
6. Thanks to the multi-connected nature of NNE measurement nodes, we can directly compare the performance and reliability of different networks at the same location, and thereby *quantify the potential gain in robustness from end-device multi-homing*. We find that there is mostly good diversity in radio conditions between operators, and that downtime can be reduced significantly if multiple networks can be used in parallel. In fact, most measurement nodes can theoretically achieve 99.999% ("five nines") connection availability when combining two operators.
7. As the last part of this thesis, we present MULTEX – a novel approach that enables mobile devices to have simultaneous access to multiple packet data networks and thus increase the reliability by built-in multi-homing feature of LTE.

1.5 Publications

Most of the work this thesis is built on has been published in conference proceedings and journals. In the following, we list all related publications.

1. A. Kvalbein, D. Baltrūnas, K. Evensen, J. Xiang, A. Elmokashfi, and S. Ferlin-Oliveira. The Nornet Edge Platform for Mobile Broadband Measurements. *Elsevier Computer Networks special issue on Future Internet Testbeds*, 61: 88–101, 2014;
2. D. Baltrūnas, A. Elmokashfi, and A. Kvalbein. Measuring the Reliability of Mobile Broadband Networks. In *Proceedings of the 2014 Conference on Internet Measurement Conference*, pages 45–58. ACM, 2014;
3. D. Baltrūnas, A. Elmokashfi, and A. Kvalbein. Dissecting Packet Loss in Mobile Broadband Networks from the Edge. In *INFOCOM*, pages 388–396. IEEE, 2015;
4. N. Larson, D. Baltrūnas, A. Kvalbein, A. Dhamdhere, C. KC, and A. Elmokashfi. Investigating Excessive Delays in Mobile Broadband Networks. In *Proceedings of the 5th Workshop on All Things Cellular: Operations, Applications and Challenges*, pages 51–56. ACM, 2015;
5. D. Baltrūnas, A. Elmokashfi, A. Kvalbein, and Ö. Alay. Investigating Packet Loss in Mobile Broadband Networks under Mobility. In *IFIP Networking Conference (IFIP Networking) and Workshops*, pages 225–233. IEEE, 2016;
6. D. Baltrūnas, A. Elmokashfi, and A. Kvalbein. MULTEX: Multiple PDN Connections in LTE and Beyond for Enhanced Routing and Services. In *Proceedings of the 5th Workshop on All Things Cellular: Operations, Applications and Challenges*, pages 13–18. ACM, 2016.

Part I

A framework and a platform for measuring mobile broadband reliability

Chapter 2

The evolution and architecture of MBB networks

2.1 The basic functions and the evolution of mobile networks

The historical and technological evolution of mobile networks is often divided into generations. Figure 2.1 shows the evolution from second generation (2G) networks towards fifth generation (5G). In the remainder of this section, we provide an overview of each generation and highlight the central design and architectural aspects. Where possible, we focus on the functions and components that are the most relevant to mobile data services.

2.1.1 2G

In 1990, the 2G cellular digital networks started to emerge, replacing earlier analogue cellular networks. The radio service of a cellular network is divided into cells, each served by at least one physical base station. Cell phones connect to a network by searching for cells in the immediate vicinity.

In Europe, Global Standard for Mobile Communications (GSM), originally called Groupe Spécial Mobile, was standardised by European Telecommunications Standards Institute (ETSI), whereas in the US, Qualcomm developed the Interim Standard 95 (IS95) standard that is based on Code Division Multiple Access (CDMA) and was also known as *cdmaOne*. GSM cellular networks provide not only digital

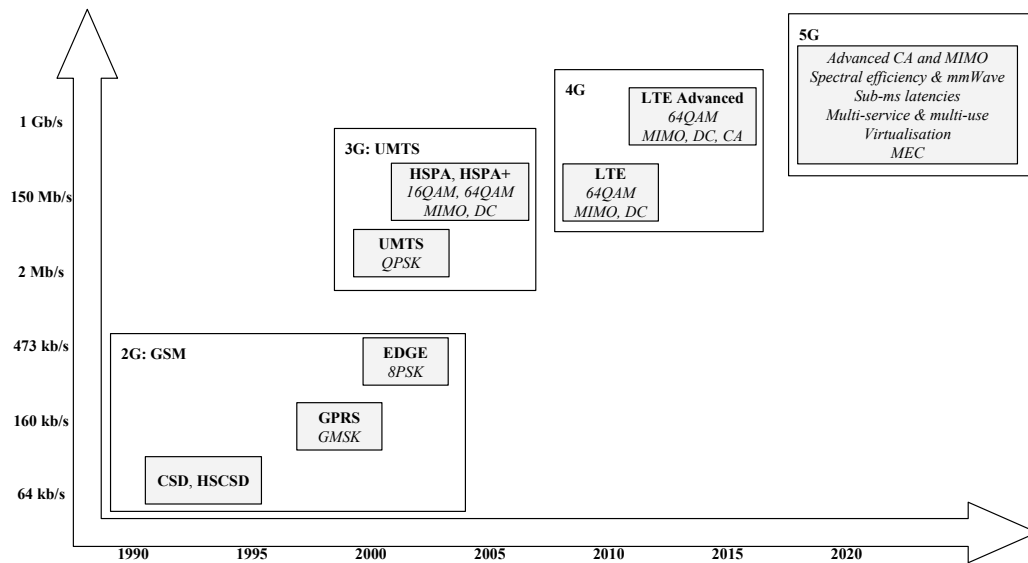


Figure 2.1: The evolution of mobile networks.

Circuit-Switched (CS) voice services, but have, when compared to analogue networks, improved capacity and coverage.

The key feature of CS services is the Quality of Service (QoS) guarantee. Certain QoS is ensured by allocating a single dedicated radio time slot between the mobile phone and the base station for the entire duration of the connection that can for example be a voice call. Another important element of a GSM network is a Subscriber Identify Module (SIM) – the smart card that is associated with every GSM user. The SIM card contains the subscriber identity and keys necessary for secure communication with the network. GSM has also introduced Short Message Service (SMS) - text messaging that is still popular today.

The architecture of GSM networks is divided into two subsystems: Base Station Subsystem (BSS) and Network and Switching Subsystem (NSS). BSS contains the base stations and their controllers, whereas NSS is responsible for switching and mobility management. NSS is often referred to as the CN. The GSM specifications define the interfaces and protocols for communication between different components. The protocols are layered according to the Open Systems Interconnection (OSI) layers. The layered protocol stack allows the software running on different BSS and NSS components to only process its own layer and then forward the packet to upper-level software or component. The exact set of protocols within each layer varies depending on the communication path.

BSS components

The BSS subsystem consists of Base Transceiver Station (BTS) and Base Station Controller (BSC). BTS contains the necessary hardware and software blocks for processing, transmitting and receiving radio signals. Several geographically close BTSs are grouped together into a Location Area (LA) that is controlled by a BSC. Typically all communication between the BTS and the BSC are encrypted. At the physical layer, the radio signals between the BTS and the User Equipment (UE) are modulated using Gaussian Minimum Shift Keying (GMSK) modulation with one bit per Radio Frequency (RF) modulated symbol. GMSK is a popular modulation scheme in radio communications due to the spectrum efficiently and relatively low interference levels.

BTS also vary in the area it covers. In rural areas, a single BTS typically covers a radius of 10 or more kilometres, whereas in cities with the high number of users there can be several BTS deployed with only a few hundred meters apart from each other.

NSS components

The NSS is controlled by the Mobile Switching Center (MSC) – the central element of a CN. MSC is responsible for call set-up, release, routing and other switching functions. It also interconnects with other Public Switched Telephone Networks (PSTNs) and Public Land Mobile Networks (PLMNs), allowing users to call or receive calls from other local or foreign mobile or fixed networks. One MSC has one or more Visitor Location Registers (VLRs) associated with it within the location or roaming area served by that MSC. The VLR is a database of subscribers that are currently roaming in the corresponding MSC's roaming area. The VLR fetches certain subscriber profile information from Home Location Register (HLR) that holds permanent subscriber data. It then caches the profile temporarily, thus minimising the number of queries between the MSC and HLR for a particular subscriber. At any point in time, a user can only be present in one VLR. For performance reasons VLR and MSC are often co-located. The VLR is also a key component to realising national or international roaming that allows users to send and receive calls and access other services from their home network while camping in a visited network. In case of roaming, a user registers on the visited network's VLR, which in turn fetches the subscriber information from the home HLR. The home HLR is updated with the visited network's VLR address, which it needs for

delivering services originating from the home network, such as mobile terminated calls and SMS messages, to the user.

CS data services

The CS data transmission in GSM is realised using Circuit Switched Data (CSD). CSD, as with Packet-Switched (PS) services, allocates a single dedicated radio slot to deliver low bitrate (9.6 kbps) data transmission to the NSS and is similar to the dial-up service used in fixed telephone providers. High Speed Circuit Switched Data (HSCSD) is considered a successor of CSD. It offers bitrates up to 57.6 kbps by exploiting more than one radio channel for the data connection. In both CSD and HSCSD, one or more allocated radio channels remain occupied for the entire duration of the connection, regardless of whether there is an ongoing data transmission or not.

GPRS

In 2000, as an integral part of GSM, General Packet Radio Service (GPRS) was introduced. GPRS is a PS data service that allows mobile terminals to access the Internet. In PS data transmission, the data is divided into small data packets and transmitted through one or more available radio channels. This way, the radio channels are only occupied when the data transmission is ongoing and are released afterwards. Therefore multiple users can access the available radio channels simultaneously. The concept of radio channel sharing makes GPRS a best-effort service, implying that latency and achieved throughput are dependent on the number of users attached to the same base station and accessing the service concurrently. This is the opposite of QoS guarantees in CS, as described above.

GPRS, as HSCSD, can use more than one time slot. When one time slot is in use, 8 kbps bitrate is available to the user. In total, 8 time slots can be used simultaneously. GPRS defines four different coding schemes within GMSK, raising the achievable user data (excluding the Radio Link Control (RLC) and Media Access Control (MAC) overhead) bitrate from 8 kbps (with *CS-1* coding scheme in use) to 20 kbps (when *CS-4* coding scheme is used). Therefore the maximum achievable bitrate for a GPRS user is 160 kbps.

GPRS requires the presence of two new elements in the NSS subsystem: Serving GPRS Support Node (SGSN) and Gateway GPRS Service Node (GGSN). The SGSN performs similar functions as the MSC for the voice traffic and is responsible for handling all PS data within the network, user authentication and mobility

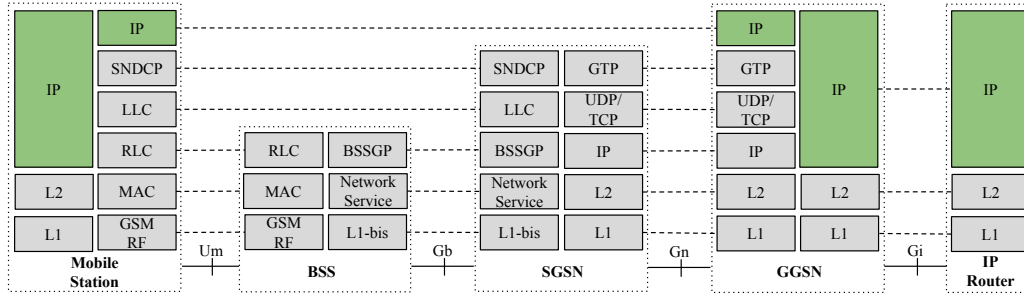


Figure 2.2: GPRS protocol stack[29, 40].

management. Similarly to the MSC, the SGSN keeps track of user location information such as current Cell Identifier (CID) and stores user profiles. A GPRS network can have several SGSNs present, deployed at different geographical locations and often co-located with MSCs. The GGSN is the gateway for the GPRS network to external networks such as the Internet and, among other functions, is responsible for Internet Protocol (IP) address allocation. SGSN and GGSN are further detailed in Section 2.2.

Figure 2.2 shows the GPRS protocol stack in the transmission plane between the Mobile Station (MS) and GGSN. One of the most important aspects to understand about the protocol stack is the fact that each GSM subsystem acts as a relay, transforming the incoming Protocol Data Units (PDUs) into the set of protocols understood by the next component. Next, we briefly detail each layer and the involved protocols.

The physical layer is divided into two sub-layers: the GSM RF layer and physical link layer. The GSM RF layer is used to control physical radio channels and modulation, and is realised by the means of Time Division Multiple Access (TDMA) and Frequency Division Multiple Access (FDMA). The physical link layer is responsible for channel coding, synchronisation and power control. Each physical channel supports a number of logical channels used for user traffic and signalling. The air interface between the MS and the BSS is called U_m .

Above the physical layer, there are three control plane layers: MAC, RLC and Logical Link Layer (LLC) layers, responsible for the logical links and their control between the MS and the BSS. The MAC layer controls access to the radio channels. The RLC layer is responsible for segmenting and reassembling the LLC PDUs into the RLC data blocks and providing retransmission mechanisms for the erroneous data blocks. The LLC layer formats the data frames and provides a

reliable ciphered link between the MS and the SGSN.

The Base Station System GPRS Protocol (BSSGP) controls the flow of LLC frames across the G_b interface by providing radio related QoS and routing for data transfer between the BSS and SGSN. The BSSGP PDUs are conveyed by the *network service* layer. The Sub Network Dependent Convergence Protocol (SNDCP) maps the IP with the underlying transport protocol stack by performing conversion, compression and segmentation of IP packets into Sub Network Protocol Data Units (SNPDUs).

The tunnelling of the user IP payload is realised by the General Packet Radio Service Tunnelling Protocol (GTP) PDUs, which are transported within the GPRS backbone network inside the User Datagram Protocol (UDP) or Transmission control protocol (TCP) datagrams. The BSS forwards the LLC PDUs between the air interface and the G_b interface, while the relay function in the SGSN forwards the Packet Data Protocol (PDP) PDUs between the G_b and G_n interfaces. The GGSN terminates the GTP tunnel and forwards the decapsulated IP packets towards an IP router using the G_i protocol.

EDGE

The 3rd Generation Partnership Project (3GPP) Release '98 introduced Enhanced Data rates for GSM Evolution (EDGE). EDGE is backwards compatible extension to GPRS and offers increased data transmission rates as well as improved network capacity. The first commercial EDGE deployment was in 2003 by Cingular Wireless in the US. EDGE uses nine Modulation and Coding Schemes (MCSs), four of which (*MCS-1* through *MCS-4*) use the GMSK modulation and the other five (*MCS-5* through *MCS-9*) employ the 8 Phase Shift Keying (8PSK) modulation[67]. 8PSK conveys three bits per RF modulated symbol, effectively increasing the bitrate to 59.2 kbps (when *MCS-9* is used) per time slot, which is almost three times higher compared to 20 kbps with GPRS and *CS-4*. When eight time slots are in use, EDGE can therefore provide a theoretical user data bitrate of 473.6 kbps.

EDGE is further extended in 3GPP Release 7 as part of the Evolved EDGE feature set. Evolved EDGE offers 1 Mbps bitrate peak throughput by using downlink dual carrier operation, modulation schemes with higher order and symbol rate, and turbo coding. Evolved EDGE also reduces latencies by lowering the Transmission Time Interval (TTI) from 20 ms to 10 ms.

2.1.2 3G

The demand for greater data bit rates and lower latencies led to the introduction of 3G networks, namely Universal Mobile Telecommunications System (UMTS) and Code Division Multiple Access (CDMA2000). The first 3G networks were deployed in 2001 by NTT DoCoMo in Japan, followed by commercial deployments in Europe in 2003. UMTS is backwards compatible with the GSM network. It uses Wideband Code Division Multiple Access (WCDMA) as the driving radio access technology with a 5 MHz channel bandwidth, offering faster data transmission on the air interface. WCDMA supports two different duplex methods, Frequency Division Duplex (FDD) and Time Division Duplex (TDD). Until the 3GPP Release 99, UMTS used Quadrature Phase Shift Keying (QPSK) modulation with two code bits per symbol, offering maximum data rates in both downlink and in uplink up to 384 kbps (up to 2 Mbps theoretical)[42]. Increased data rates enabled users residing at locations with poor fixed broadband coverage to switch to 3G networks as their primary Internet provider.

The network protocol stack and the CN network architecture of UMTS networks is relatively similar to that of the GSM networks. In the following, we list the major changes with respect to the terminology and the changes in architecture:

1. The GSM MS is referred to as UMTS UE.
2. The GSM BSS that constitutes the Radio Access Network (RAN) in GSM networks is referred to as Radio Network Subsystem (RNS) of the UMTS Terrestrial Radio Access Network (UTRAN).
3. The GSM BSS components, BTS and BSC, are referred to as the UTRAN UMTS Node B (NodeB) and the UTRAN RNC, respectively.
4. New states are introduced in the Radio Resource Control (RRC) state machine (see Section 2.2.2).
5. In UMTS, SNDCP is replaced by Packet Data Convergence Protocol (PDCP), which is responsible for transferring and compression of user and control plane data to the RRC and upper layers.
6. The video calling feature of the UMTS networks requires a set of IP Multimedia Subsystem (IMS) components, such as Media Gateway (MGW), to be present within the CN.

7. In UMTS networks, there can be a number of optional components for real-time data charging and other value added services, such as video calls.

Later 3GPP releases introduced Hight Speed Packet Access (HSPA). HSPA is used to refer to either Hight Speed Downlink Packet Access (HSDPA), Hight Speed Uplink Packet Access (HSUPA) or Evolved Hight Speed Packet Access (HSPA+), which are based on the Quadrature amplitude modulation (QAM) modulation scheme with four (16-QAM), five (32-QAM) and six (64-QAM) code bits per symbol, respectively. HSPA significantly increases the user data bitrate that, when combined with Dual Carrier (DC) and Multiple-Input Multiple-Output (MIMO), can exceed 168 Mbps in the downlink and 22 Mbps in the uplink.

In 3G, the concept of small cells, deployed indoors in larger corporate offices and interconnected with the NSS over the existing high-speed Internet connection available at the premises, has been developed and became popular across all generation MBB networks.

2.1.3 4G

4G is the latest generation commercial MBB networks, succeeding the former 3G. 4G aims at further reducing latencies, improving the data rates and spectral efficiency. There are two major 4G network technologies: LTE and Worldwide Interoperability for Microwave Access (WiMAX). The first 4G network was proposed by NTT DoCoMo in Japan in 2004 and in 2008 standardised in 3GPP Release 8. The first publicly available LTE network was launched by TeliaSonera in Oslo and Stockholm in December 2009. WiMAX is based on the Institute of Electrical and Electronics Engineers (IEEE) 802.16 series of standards for wireless metropolitan area networks. Hereafter, we will only focus on LTE, currently the most widely deployed 4G MBB network.

In contrast to the previous generation networks, LTE is an all-IP network and introduces numerous new features and enhancements. LTE uses Orthogonal Frequency-Division Multiple Access (OFDMA) for the downlink and single-carrier FDMA for the uplink and supports two different duplex methods, FDD and TDD, formally referred to as LTE-TDD and LTE-FDD[43]. OFDMA modulation, when compared to WCDMA and CDMA, has, among other features, higher spectral efficiency and is more robust to channel interference. The OFDMA signal is divided into 2048 sub-carriers with 15 kHz spacing and modulated using QAM with 2, 4, or 6 bits per symbol. In Europe, almost all networks use LTE-FDD, whereas

LTE-TDD is more popular in the rest of the world. LTE spectrum is divided into a number of LTE-FDD and LTE-TDD bands with the allocated frequencies for the uplink and downlink. LTE supports scalable channel bandwidths, ranging from 1.4 MHz to 20 MHz.

Higher available bandwidth and the use of massive MIMO and Carrier Aggregation (CA) techniques specified in 3GPP Release 10 (also referred to as LTE Advanced) lead to increased user data rates. The TTI in LTE is reduced to 1 ms that in turn reduces the overall end-to-end latency for the user data payload. Depending on how many antennas for the downlink MIMO it has and what modulation schemes it supports, an LTE UE can be assigned to one of the 11 LTE categories. Most modern LTE UEs today belong to LTE Category 6, which, when the 20 MHz channel bandwidth is used, provides up to 301.5 Mbps and 51 Mbps theoretical physical layer data rates in the downlink and uplink direction, respectively.

LTE introduces several major changes to the network architecture:

1. E-UTRAN Node B (eNB), an LTE base station uses IP to communicate with the CN, replacing the BSSGP used in GSM and UMTS and making LTE an all-ip network.
2. eNB performs the RRC control and other functions that were previously governed by RNC or BSC.
3. LTE provides only PS services; voice or video calls are implemented using Voice over LTE (VoLTE) or Circuit Switched FallBack (CSFB) - a technique that forces UEs to downgrade to UMTS for the duration of the call.

With respect to the network protocol stack and the CN architecture, LTE reuses most of the concepts from the previous generation networks. User data payload is encapsulated into the GTP tunnel towards the centralised Evolved Packet Core (EPC).

2.1.4 5G

The next-generation MBB network, often called 5G, features a number of improvements and architectural changes to the traditional MBB architecture. 5G promises sub-millisecond latencies, data rates in the order of gigabits per second and the greater diversity in terms of the radio spectrum and available RATs.

One of the main trends of 5G is that most radio access and core network infrastructures are expected to be cloud-based, enabling faster scalability and smoother

network upgrades. The majority of 5G CN components are seen to be run as virtual machines in geographically distributed clouds, interconnected together by the means of Software-Defined Networking (SDN). The concept of Network Function Virtualisation (NFV) is specified by ETSI and is being adopted to pilot 5G networks worldwide.

5G also proposes different approaches for virtualising RAN components. The main concept of a virtualised RAN is that the functions of a base station are split into two parts. The radio function unit, Remote Radio Head (RRH) is connected by the fast medium such as fibre or microwave antennas to the Baseband Unit (BBU) that resides in the cloud and can be organised in pools if necessary. Such concept not only potentially reduces latencies, but also gives network operators more flexibility in managing and planning the network resources as well as adding support for the latest radio access technologies without the need to physically upgrade the base stations.

Another prominent development path in 5G is Mobile Edge Computing (MEC) that is also being specified by ETSI. The main idea behind MEC is to bring virtualised computational resources closer to the edge, enabling new types of applications, reducing network congestion and off-loading the CN. MEC components are seen to be deployed within the base station such as eNB. MEC also has a notion of content storage and processing that creates potential for realising content caching at the edge.

2.2 MBB network architecture

Measurements conducted and analysed in this work are from 2G, 3G and 4G MBB networks. More specifically, in this work we studied and collected the data from GSM, UMTS, CDMA2000, and LTE operational networks. The remainder of this chapter describes the more detailed architecture of MBB networks, focusing on the PS services.

As described in Section 2.1, although GSM, UMTS and LTE networks belong to different generations, their conceptual architecture is similar. Figure 2.3 shows the main components of GSM and UMTS networks, whereas Figure 2.4 depicts the simplified architecture of an LTE network. All three networks are logically divided into the RAN and the CN. In UMTS, RAN is referred to as UTRAN, whereas in LTE it is referred to as Evolved UTRAN (E-UTRAN). Further, in LTE the CN is referred to as EPC. In the remainder of this thesis, we use *RAN* and *CN* to refer to

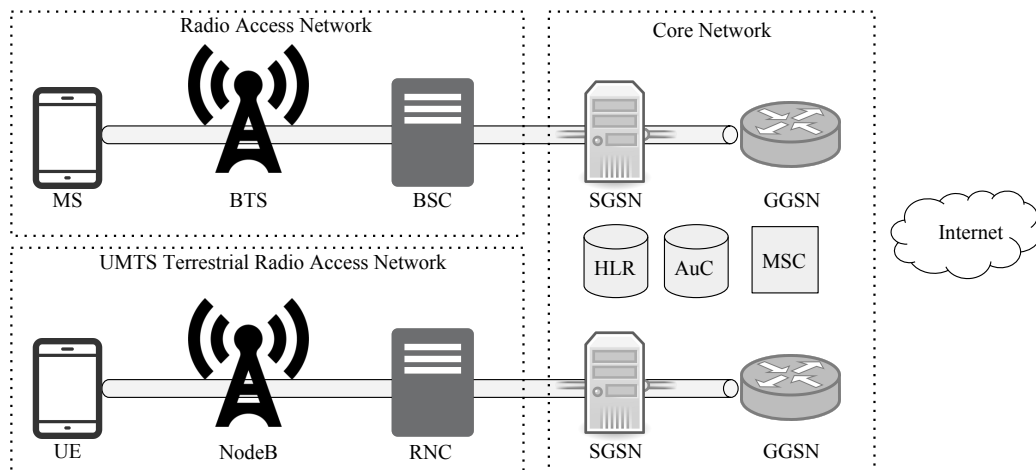


Figure 2.3: Simplified GSM/UMTS MBB network architecture.

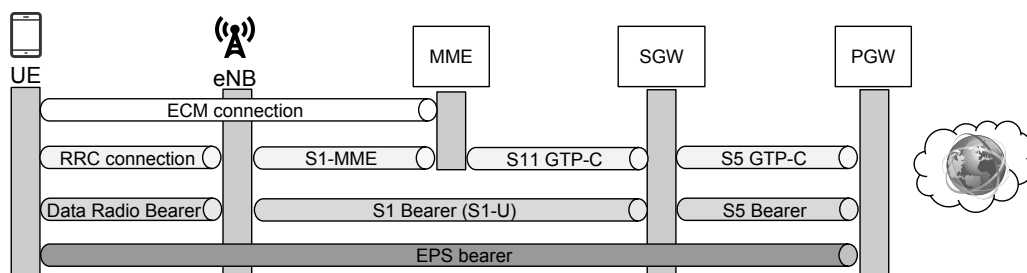


Figure 2.4: Simplified LTE network architecture.

the radio access or the core networks, respectively, regardless of the RAT.

The RAN resides between a mobile device such as smartphone or modem, and the CN. The CN consists of centralised databases and registers to store subscriber information as well as gateways and supporting systems that are necessary to provide CS and PS services to the users. In the following, we describe the functions of the RAN and the CN in more details.

2.2.1 RAN

The RAN of a MBB network consists base stations, and, in case of GSM and UMTS, base station controllers. Depending on the RAT, different terminology is used. In GSM, mobile device is called GSM MS, and in UMTS and LTE networks, it is referred to a UE. The GSM base station and its controller is referred to as BTS and BSC, respectively. In UTRAN, the base station is referred to as NodeB, and the base station controller is called RNC. In E-UTRAN, the base station is referred to as eNB, where it also replaces RNC for the radio resource management.

A set of GSM or UMTS base stations are grouped geographically into LA, identified by LAC. Similarly, LTE eNBs are grouped into multiple Tracking Areas (TAs). Typically, but not necessarily, one BSC or RNC controls and hence corresponds to one LAC, whereas one LTE TA is served by one Serving Gateway (SGW) that is part of the EPC. The connectivity between the base stations and their respective controller, often referred to as back-haul is typically realised by the means of fibre or microwave transmission.

A base station is typically divided into several sectors. Each sector is identified by the CID and has a locally unique radio frequency channel allocated to it. Since one cell serves multiple users simultaneously, mechanisms for spectrum sharing and resource management must be in place. Most of radio resource management functions are governed by a BSC/RNC controller or an eNB. The controller or the eNB allocates necessary radio resources to the users. In case of GSM or UMTS networks, the controller also implements the mobility management functions that enable users to handover between different base stations.

2.2.2 RRC state machine

Before a MS (UE) is allowed to communicate with the CN components or the Internet, it first needs to acquire a Radio Access Bearer (RAB). Radio resource management is realised by the means of the RRC protocol. The RRC defines a

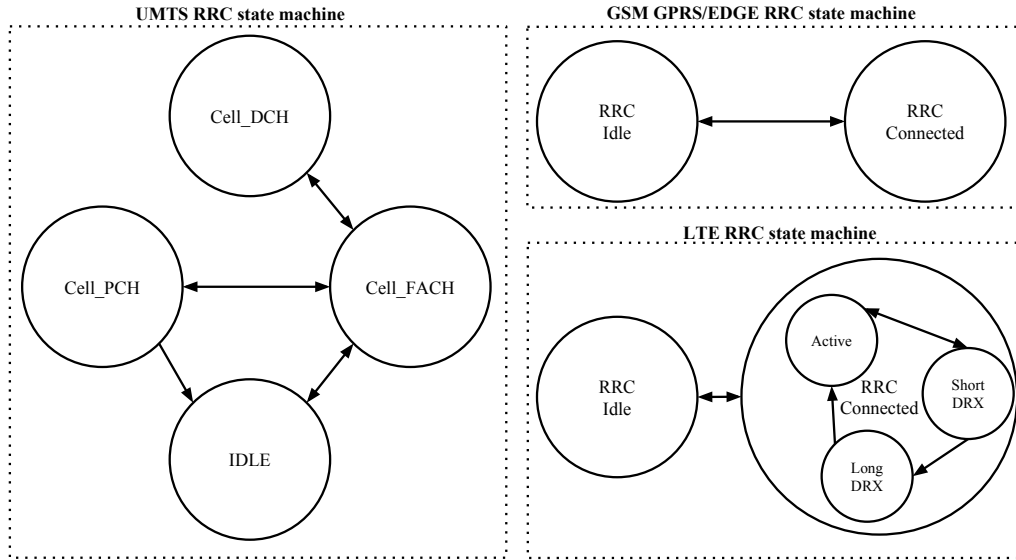


Figure 2.5: RRC state machines of GSM, UMTS and LTE networks.

state machine that include certain states that a MS (UE) may be in. The RRC state of a user is controlled by GSM BSC, UMTS RNC, or LTE eNB and is dependent on the traffic pattern, which is typically a combination of a bitrate and a set of idle timers. Different RRC states are associated with either shared or dedicated radio channels and respective RABs.

Figure 2.5 shows the RRC state machines of GSM, UMTS and LTE networks. If the user is not sending any data for a preconfigured amount of time (typically in the order of seconds), RRC demotes the state to *IDLE* (or to *CELL_PCH* in some UMTS networks). Otherwise, based on the bit rate, a UMTS user can be promoted with a *CELL_FACH* state (shared channel, low bit rate, low power usage) or a *CELL_DCH* state (dedicated channel, high bit rate, high power usage), whereas a GSM or LTE user is transited to a *RRC CONNECTED* state.

While the RRC state machine helps to save radio resources in terms of shared and dedicated radio channels, it creates extra complexity and increases end-to-end delays due to different timers of the state machines and transition between states. Moreover, frequent state transitions drains UE's battery. To deal with the side effects of frequent RRC state transitions, some HSPA and LTE networks implement Discontinuous Reception (DRX) mechanisms, defined in 3GPP Releases 6 and 7. Even when a UE is on *IDLE* RRC state, it still needs to constantly listen to the paging messages on the physical downlink channel. DRX introduce preconfigured

long and short sleep cycles, during which a UE does not need to monitor paging messages and hence can increase battery life.

2.2.3 PDP context, EPS bearer and data tunnelling

GSM and UMTS networks provide CS and PS services to the users. Unlike previous generations of mobile networks, LTE is an all-IP network, meaning that all signalling and data between the network components is exchanged over IP. A fundamental design aspect of all generation mobile networks is the concept of tunnelling data in the PS domain sent from or destined to a MS/UE all the way to the CN. The concept of tunnelling makes mobile networks conceptually different from the rest of the Internet, where the route for an IPflow involves a set of geographically distributed IProuters, each making an autonomous routing decision on where the IP packet has to be forwarded further.

In the PS domain, all IP packets sent by a MS/UE are encapsulated in a GTP tunnel that is associated with an established PDP context (in GSM and UMTS) or an Evolved Packet System (EPS) bearer (in LTE). The PDP context (EPS bearer) is a data structure that contains the IP address and other information about the user session.

A single GTP tunnel (and its respective PDP context or EPS bearer) is associated with a single Access Point Name (APN) that provides access to one Packet Data Network (PDN). In other words, a MS (UE) must have an established PDP context or an EPS bearer as well as a non-idle RRC state (as described in Section 2.2.2) in order to be able to exchange the IP packets with a PDN, such as the Internet.

An LTE UE that is attached to a network must also have a default EPS bearer established at all times. The default EPS bearer can have one or more dedicated EPS bearers associated with it and each EPS bearer is assigned with different Quality of Service Class Identifiers (QCI). Further, the two bearer types differ in a way that the dedicated EPS bearer can get a guaranteed bit rate assignment, whereas the default EPS bearer is always best-effort.

2.2.4 CN components

In GSM and UMTS networks, GTP tunnel endpoints are the BSC/RNC and a GGSN, whereas in LTE networks, the endpoints are the eNB and a Packet Data Network Gateway (PGW). The GGSN (PGW in LTE) is the central component

of the GPRS network and it is the first IP hop towards any host in the Internet. It is responsible for user authentication, session management, IP address allocation, IP level security, routing the IP packets to specific PDNs or the Internet. In GSM and UMTS, the GTP is further split into GTP user plane (GTP-U) between BSC/RNC and SGSN and GTP control plane (GTP-C) between SGSN and GGSN. In LTE, the functions of SGSN are handled by the SGW and Mobility Management Entity (MME). SGSN (SGW in LTE) is the packet data service termination point towards RAN and delivers data packets to and from mobile users that reside in SGSN's (SGW's) serving area. SGSN is also responsible for authentication and mobility management, whereas in LTE, all control plane functions for subscriber, session and mobility management are handled by MME.

SGSN and GGSN are the central CN components for PS services in GSM and UMTS networks, and their respective counterparts in LTE are SGW, PGW and MME. SGSNs (SGWs) are geographically distributed to serve users from a particular area. Multiple SGSNs (SGWs) are linked to one or more GGSN (PGW). Depending on the size of a MBB network, IP traffic demands, and reliability requirements, multiple GGSNs (PGWs) can be present within the network and are either located in the centralised CN or distributed across an operator's data centres.

In addition to the aforementioned components for the PS domain, the CN includes a set of other elements to realise various network functions. The central element of a CS domain in GSM and UMTS networks is MSC that controls the network switching subsystem elements. Subscribers along with their profile and authentication data is stored in HLR and Authentication Centre (AuC), which in LTE networks are replaced by Home Subscriber Server (HSS). LTE has introduced VoIP services, realised by IMS components within the CN. GGSN and PGW can optionally be connected to the Policy and Charging Rules Function (PCRF) component that defines policies for real-time charging. Operators also put significant efforts to ensure better QoS and QoE by employing different optimisation mechanisms realised by middleboxes such as deep packet inspection nodes or TCP proxies that typically reside in the CN.

Chapter 3

Measuring MBB network reliability

3.1 The definition and importance of reliability

Historically, the major criteria differentiating MBB providers have been price, data throughput and coverage. However, the growing number of critical services, such as electronic ticketing system in public transport and credit card payment terminals, that depend on MBB as a bearer for sensitive transactions, demand *reliable* and *resilient* network connectivity.

In the literature, there exist several definitions of *resilience* and *reliability*. In the context of *ResumeNet* project and the *ResiliNets* initiative, resilience is defined as the ability of the network to provide and maintain an acceptable level of service in the face of various faults and challenges to normal operation [92]. The authors also define the two integral parts of resilience: challenge tolerance and trustworthiness. They assign availability and reliability as part of trustworthiness and denote reliability as the continuity of service, that is the probability that a system or service remains operational for a specified period of time. The EU project *METIS* defines reliability as an assessment criterion to describe the quality of a radio link connection for fulfilling a certain service level [45]. The authors term reliability as the probability that a certain amount of data to or from an end user device is successfully transmitted to another peer within a predefined time frame that corresponds to the service requirements, such as end-to-end latency.

Reliability is an important contributor towards better user experience and reliable MBB services. According to a survey of broadband users conducted in 2011, reliability is one of the main sources of dissatisfaction with broadband service [61]. In order to state whether a MBB network is reliable or not, we need to measure reliability in a way that provides objective outputs to evaluate.

3.2 Reliability metrics

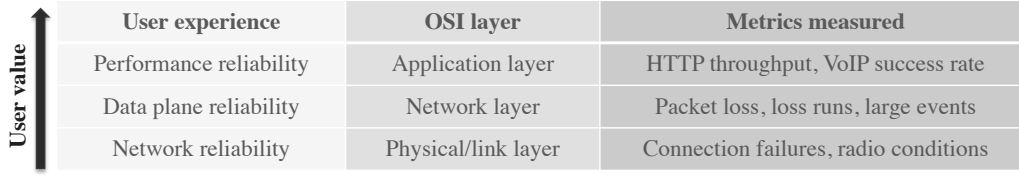
Several ongoing efforts are aiming to standardise the task of performing large-scale measurements. The Internet Engineering Task Force (IETF) IP Performance Metrics (IPPM) Working Group is developing metrics for the quality, performance, and reliability of Internet data delivery services and applications. In 2012, the IEEE 802.16 Working Group on Broadband Wireless Access Standards initiated the *Mobile Broadband Network Performance Measurements* project to define procedures for characterising the performance of deployed mobile broadband networks from a user perspective [46]. The project has however been terminated in 2016. Bagnulo et al. [10] defined meaningful performance metrics for IP networks and showed that the actual tests performed by existing platforms are loosely related to existing standards, as the existing standards are simply too generic to be useful in an operational platform.

In the studies presented in this thesis, we use simple and widely accepted metrics to measure reliability. We assess the MBB network reliability at three different levels: control, data, and application. We use standard tools and Internet protocols to continuously monitor the health of MBB connections and periodically generate traffic over them.

3.3 Proposed framework for measuring reliability

We believe that reliability in MBB networks is too complex to be understood only through static analysis of the components involved, and that the most promising approach for assessing and predicting the reliability of the offered service is through long-term end-to-end measurements. We argue that reliability must be characterised at several levels, including the basic connection between the UE and the base station, the stability of the data plane, and the reliability of application level performance. In particular, we are interested to know a) how long a network connection will survive, b) how durable and frequent are the downtimes, c) how frequently the data service is unavailable, d) how reliable is the data service, and e) how stable is the application application performance. In this work, these aspects of reliability are assessed through long-term active measurements from a large number of geographically distributed measurement nodes that are part of the NNE platform, described in Chapter 4.

Reliability relates to several stability and performance related metrics. In order to measure these different types of notions of a particular MBB network, it is



User experience	OSI layer	Metrics measured
Performance reliability	Application layer	HTTP throughput, VoIP success rate
Data plane reliability	Network layer	Packet loss, loss runs, large events
Network reliability	Physical/link layer	Connection failures, radio conditions

Figure 3.1: Framework for measuring experienced reliability in MBB networks.

necessary to decide on the type of measurement tests and metrics that would let to make certain statements about network reliability. Here, we propose a model where the reliability of a network is measured at different levels, reflecting increasing value for the user. A high level picture of the framework is shown in Figure 3.1.

The proposed model is a generic framework for describing the experienced reliability in MBB networks. We want to test the stability of an established network connection and the overall availability of the network. By looking at measurements from individual connections, we are able to identify important differences between networks and to characterise the reliability of each network as a whole. In this work, we select a few relevant metrics at each level, and use these to characterise reliability of the measured networks. We check if it is possible to establish a MBB connection and measure the time between failures by monitoring the state of a PDP context (EPS bearer). During the lifetime of an active connection, we assess data plane reliability by continuously sending small UDP packets to an echo server and recording round-trip time (RTT) and packet loss. At the application level, we use well known application-level protocols and analyse the impact of the underlying levels on the application level performance. Other metrics can later be added to give an even more complete picture.

In the following, we describe the three levels in more details.

3.3.1 Connection level reliability

At the very basic level, the UE should have a reliable connection to the MBB network. By "connection" in this context, we mean that there is an established PDP context (EPS bearer) in the CN. The establishment of a PDP context (EPS bearer) is a complex operation by itself, involving negotiation between the UE and different CN components. The involvement of many different network components increases the chance that authentication or authorisation might not be successful,

leading to the failed establishment attempts. The stability of the PDP context depends on both the RAN and the CN; the PDP context (EPS bearer) can be broken by loss of coverage, failures in base stations or transmission, or by failures or capacity problems in the central components such as SGSN (SGW) or GGSN (PGW).

From the UE side, having a PDP context (EPS bearer) maps to having an assigned IP address from the mobile network. In Section 6.1, we measure reliability at the connection level by looking at the stability of the IP address assignment as a proxy for the PDP context using the stationary connections. The metrics we look at are how often the connection is lost, and how long it takes before the PDP context can be successfully re-established. We also analyse how these metrics are related to underlying characteristics of the connections, such as signal strength and connection RAT. The selected metric describes the stability of connections over time.

3.3.2 Data plane reliability

Having an established PDP context (EPS bearer) does not necessarily mean that the UE has well-functioning end-to-end connectivity to the Internet. MBB networks tend to have a significantly higher latency and jitter when compared to the wired ones, resulting in longer connection establishment times and responses from the end-host. For the user, it means longer waiting times for a response after making a network request (for example, to open a web page). Real-time applications such as gaming or VoIP calls fall into delay-intolerant application traffic [60], which means that higher latency will directly impact user experience. Reduction of delay is one of the main goals for next generation wireless networks [58]. Interference, drop in signal quality or congestion in either the wireless access or elsewhere in the mobile network may disrupt packet forwarding. This can not only cause periods of increased delays, but also result in excessive packet loss, or "gaps" where no data comes through.

In Sections 6.2, 6.3 and 7.1, we measure data plane reliability by looking at loss patterns in long-lasting continuous probing streams from our large-scale measurement study. We demonstrate that end-to-end active measurements can give invaluable insights into the nature and characteristics of packet loss in cellular networks. We first describe loss patterns in each network, and discuss how loss must be seen in relation with the radio condition of the MBB connection. We use packet loss to identify abnormal events where packet loss is higher than normal for a significant number of connections. We then investigate the possible causes of loss and find

that a significant fraction of loss occurs during pathological and normal RRC state transitions. We also observe loss with pronounced diurnal patterns and relatively strong correlation between geographically diverse measurement locations. Some of our results indicate that the causes of a significant part of loss lie beyond the RAN.

In Sections 6.4 and 7.2, we report the results of a measurement study of round trip delays. We find high variation in delay within the same RAT and RRC state. We also observe connections with round trip delays of several seconds, often multiple times per hour, that are more frequent when the connections are under mobility. We correlate these extreme delay events with available connection metadata and find that they are related to handovers, RRC state transitions, and retransmissions at the link and physical layers.

3.3.3 Application level reliability

Reliability also involves a notion of stability and predictability in the performance an application achieves over the MBB network. This stability depends of course on both the connection level reliability and the data plane reliability. Application layer performance varies depending on the specific application requirements. Some applications will perform well under a wide range of network conditions, while others have stricter requirements on available bandwidth or delay. In MBB networks, the experienced network performance depends on the *state* of the connection, since radio resources are assigned depending on the traffic load. It is therefore difficult to predict the performance of an application based on generic measurement probes. Instead, application performance should be assessed through experiments with actual application traffic.

In Section 6.5, we look at the stability in performance of two popular applications using the stationary MBB connections and find that pauses in packet forwarding are likely caused by the lack of available radio resources.

3.4 Related approaches for measuring MBB networks

During the past few years, there have been a growing interest by regulators, policy makers and the networking community in measuring the performance of MBB networks. Regulators need measurements to monitor how operators fulfil their obligations, and as a baseline for designing regulatory policies. On the other hand, operators themselves are interested in operational instability and anomalies to identify

problems in their networks. Operators or independent agencies sometimes perform drive-by tests to identify coverage holes or performance problems. These tests are, however, expensive and do not scale well. The challenges for measuring MBB performance are very different from the fixed-line networks. Many different factors, such as signal strength, device model, RAT, frequency of handovers and location information of mobile devices need to be taken into account when conducting MBB measurements [11]. Another challenge with end-to-end mobile measurement tools is that they demand incentivising a statistically significant sample of users to install and run the tool as well as protecting those users' resources from abuse and preserving user privacy [31]. Users typically cannot measure their MBB provider's performance themselves due to restrictive data plans and usage caps. Mobile devices are not homogeneous in terms of the variety of mobile operating platforms and can often show different performance measurements even for the same MBB network. Such diversity is especially challenging for mobile application developers, who can only evaluate their applications on a limited number of devices and networks.

Several regulators have translated the aforementioned challenges into ongoing nationwide efforts to conduct long long-term continuous monitoring of MBB networks. Federal Communications Commission (FCC), the national regulator in the United States, initiated the *Measuring Broadband America* [32], a nationwide program to study the performance of broadband service, including MBB. A study by Sundaresan et al. [94], partly based on the FCC data, investigated the suitability of different metrics in characterising home broadband performance and demonstrated that the long-term continuous monitoring from the edge is indispensable for understanding and assessing the performance. The Communications Regulatory Authority of Lithuania is monitoring the quality of wireless Internet access services since 2011 [26]. They developed a monitoring system that regularly performs controlled measurements to allow users to objectively evaluate the quality of Internet access services provided by Lithuanian mobile communication operators. Several performance metrics are collected, including connection availability, *HyperText Transfer Protocol (HTTP) GET*, and File Transfer Protocol (FTP) throughput. The measurement results are published on a regularly updated, interactive and publicly available map.

Faggiani et al. [31] classifies end-to-end mobile network measurement, monitoring, and experimentation platforms into three categories: research testbeds for network experimentation, extensible distributed measurement tools, and services for widespread monitoring of networks performance. With respect to the way

measurement data is collected, there are several possible approaches to performing systematic measurements of MBB reliability and performance. One way is to use dedicated measurement nodes and a backend infrastructure for continuous active or passive measurements of MBB networks. Another approach is to crowd-source measurement results from large number of MBB users to get insights into different aspects of MBB reliability. MBB performance can also be assessed using network-side data, which can help in explaining the bottlenecks and degradations that cannot be detected from the UE.

Our approach is to use a dedicated measurement platform and perform continuous end-to-end measurements to measure the MBB as experienced by the user. This thesis is the first to present a country-wide assessment of MBB reliability filling an important gap in this area. In the remainder of this section, we provide an overview of the existing approaches for measuring reliability of MBB networks and contrast them to our chosen approach.

3.4.1 Platforms with dedicated measurement nodes

Dedicated measurement infrastructures are used by several projects and initiatives to conduct continuous measurements of home and mobile broadband networks. Dedicated measurement nodes provide greater flexibility with respect to the possible experiments and metadata when compared to for example crowd-sourced approach. In the following, we provide an overview of the past and current examples of such approaches. Only one of the infrastructures with dedicated nodes focuses on MBB networks, while the other projects are for measuring home broadband performance.

SamKnows

In 2009, a UK based company *SamKnows* was appointed by Ofcom, the UK telecoms regulator, to carry out the first national survey of home broadband performance. The main goal of these tests is to see if consumers receive the capacity that they pay for, but the tests also collect several other performance and functionality metrics [84]. *SamKnows* produces small, dedicated hardware probes, called *whiteboxes* that are transparent to the users, and provides measurement infrastructure. Users acting as hosts in return get monthly reports about their broadband performance, that includes end-to-end latency, packet loss and jitter, uplink and downlink throughput and goodput, and other tests. As of 2016, *SamKnows* con-

duct home and mobile broadband measurements in 30 countries, mainly ordered by telecom regulators. Their mobile *whitebox* runs a distribution of Linux, derived from OpenWrt, and has one or more Universal Serial Bus (USB) MBB dongles attached to it. The exact set of measurements is not documented, but the source code is available [83] with latest updates from 2013.

BISmark

A similar to the *SamKnows* measurement platform [94] approach is taken by BISmark project – an initiative run by Georgia Tech for measuring broadband performance. BISmark consists of several hundred dedicated, OpenWrt-based probes with off-the-shelf hardware that are deployed at volunteers' premises around the globe and support both active and passive measurements. BISmark measures, among other, end-to-end latency, jitter and packet loss, uplink and downlink throughput and against their measurement infrastructure, which mainly consists of M-Lab servers hosted by Google.

RIPE Atlas

RIPE Atlas is a measurement infrastructure operated by RIPE [90]. RIPE Atlas has several thousand hardware probes that are OpenWrt-based and deployed at volunteers' premises all over the world. RIPE Atlas probes perform only active measurements, realised by sending Internet Control Message Protocol (ICMP), *traceroute*, HTTP and Secure Socket Layer (SSL) packets to their measurement infrastructure. To deal with relatively large set of measurement nodes, RIPE Atlas has quite complex backend architecture that consists of a registration server and a set of controllers, each associated with are less than 500 probes.

WiScape

University of Wisconsin-Madison conducted a MBB performance study that was based on the WiScape [85] testbed. WiScape consisted of a handful of nodes mounted on public buses, each equipped with MBB connections from multiple operators. WiScape infrastructure was used to characterise and compare the performance of these networks in a metropolitan area in the US. Network layer performance was estimated by measuring TCP and UDP throughput, RTT, UDP loss, and application-level jitter.

Our work falls into the category of measurements using a dedicated infrastructure. In contrast to most of the related work, it is based on a much larger deployment in terms of the number of measured operators, the number of connections, geographical distribution and duration. In 4, we present a large-scale tested that consists of hundreds of dedicated measurement nodes and a backend infrastructure, operational in Norway since 2013. Further, unlike some of the related initiatives, we perform our measurements in an end-to-end fashion. In some cases, we combined end-to-end measurements collected using a dedicated infrastructure and network-side logs to provide better explanations for the observed behaviour.

3.4.2 Crowd-sourcing platforms and applications

By distributing dedicated measurement software in the form of easily installable mobile applications, researchers can harvest millions of performance tests from a large user base. In the following, we provide an overview of the most notable crowd-sourcing platforms and applications.

MobiPerf

MobiPerf [4] is an Android application that measure several performance metrics including throughput and latency and also aims to diagnose problems with application content delivery on mobile devices. MobiPerf was developed as a collaboration of University of Michigan, Northeastern University, University of Washington, and Google’s M-Lab. The measurements performed by Mobiperf are run against M-Lab servers and anonymised results are collected for research purposes. MobiPerf measures latency using ICMP if available, otherwise it falls back to measuring latency using TCP three-way handshake in a HTTP transfer. Throughput is measured by exchanging random data with a nearby M-Lab server for 16 seconds, followed by the computation of uplink and downlink throughput from packet traces. MobiPerf also supports measuring packet loss and one-way delay using UDP and can also use RRC state to estimate packet latency [81].

Nikraves et al. [70] used an extensive data set contributed by the *Mobiperf* and *Speedometer* users, and performed a longitudinal analysis of MBB performance. The authors highlighted significant performance differences across operators, access technologies, and regions. The authors also showed that the performance of MBB is not improving overtime, which necessitates the continuous monitoring of MBB performance.

OpenSignal

OpenSignal [72] is an application for smartphones that collects information about MBB coverage and throughput from the users that have the application installed. The collected data is aggregated and displayed publicly, letting users to compare the performance of different MBB networks worldwide. OpenSignal application uses both active and passive measurements for measuring latency and throughput. Both latency and throughput is measured using HTTP. In addition to measurement results, OpenSignal also records various connection metadata such as RAT, signal strength, location information and other. This information can be accessed using their public Application Programming Interface (API).

Ookla Speedtest mobile applications

Ookla's speedtest.net is one of the most widely used tools to measure the throughput. Since 2009, Ookla released the mobile version of the tool for different smartphone platforms [71]. The mobile version of the Speedtest application selects the closest measurement server based on the device location. It then uses TCP with fallback to HTTP to test the download and upload speed, latency and jitter. In case of TCP, the client establishes multiple TCP connections towards the server and adjusts the chunk and buffer size based on the real-time speed of the transfers. In case of HTTP, small binary files are uploaded (downloaded) to estimate the upload (download) speed and to select the fixed size file that is uploaded (downloaded) multiple times.

Sommers and Barford [88] used crowd-sourced data from *Speedtest.net* to compare the performance of MBB to WiFi. The authors record faster download speeds and lower latencies when using WiFi and also find low performance consistency for wireless access networks in general when compare to the fixed ones.

Netradar

Netradar is a mobile measurement platform run by Aalto University [89]. It provides tools to measure a wide-range of key network performance indicators and to automate the reasoning of the measurement results. Netradar collects measurements from their mobile application, installed by thousands of volunteers across the globe. They measure RTTs and TCP throughput in the downlink and uplink against their backend infrastructure and also passively collect metadata, such as Received Signal Strength Indicator (RSSI), RAT, and location information. Netradar data was

used to show that signal strength has low correlation to TCP goodput and that the bandwidth is degraded severely due to poor provisioning and congestion at the base station.

Portolan

Portolan is crowd-sourced mobile measurement platform operated by the University of Pisa and the Informatics and the Telematics Institute of the Italian National Research Council [31]. Portolan contains a backend servers and proxies and the smartphone application that supports active and passive measurements for signal strength, latency, forwarding path with respect to autonomous systems in the Internet, and achievable bandwidth.

myspeedtest

myspeedtest [66] is an application for Android smartphones, developed by the Network Operations and Internet Security Lab at Princeton University. The application measures downlink and uplink speed in 40 seconds and also performs latency and traceroute measurements in real-time. In addition to the speed test, it also monitors application traffic data. *myspeedtest* application is also integrated with the BISmark project.

SamKnows smartphone application

Part of *SamKnows* MBB measurements are based on an application for Android and iPhone, which is available in the US, Hong Kong and Brazil. Their smartphone application was originally designed as the *FCC Speed Test* and is part of the aforementioned *Measuring Broadband America* program. The application regularly collects performance data from volunteers and measures download and upload speed, latency and packet loss.

Speedometer

Between year 2011 and 2013, several thousand volunteers, mainly Google employees, measured MBB performance using a *Speedometer* for Android [37]. The collected data includes ICMP ping, traceroute, Domain Name System (DNS) lookup, HTTP fetches, and UDP packet loss measurements.

MITATE

Another example is MITATE [36] – a collaborative platform for mobile network experimentation, developed at Montana State University. MITATE depends on leveraging participating volunteers’ devices for running experiments requested by other participants in a tit for tat fashion. Experiments execute on user-volunteered devices that meet measurement criteria, such as MBB provider, battery level or signal strength. Participating devices periodically poll experiments from a central server. Experiments include RTT, one-way delay, throughput, jitter and loss. Measurement results are stored in a database that participating users can get access to.

PhoneLab

PhoneLab is a programmable smartphone testbed run by University at Buffalo [68] since 2013. PhoneLab experiments are performed on rooted Android devices that are given to student volunteers at the University at Buffalo. It was considered as the only testbed providing researchers with the ability to study and experiment with the Android platform. PhoneLab applications can run experiments in the background or interactively and also support Operating System (OS) level experiments. As of 2016, the testbed does not seem to be active anymore.

One of the main advantages with crowd-sourced measurements is the ability to collect data from a large variety of smartphones and MBB networks. However, we find it difficult to rely on user-initiated measurements for the continuous monitoring of MBB networks’ stability. While this approach is scalable in a way that it can collect millions of measurements from different regions, MBB networks and UEs, such measurements can be biased by the glitches and interrupts on the UE, caused by the reasons not related to the MBB network. It is also difficult to collect data on stability and availability with user-initiated measurements, since this typically requires long and uninterrupted measurement sessions. Further, users perform measurements at their own will, which might give a bias towards situations when the user experiences performance problems. Another challenge with this approach is privacy: data collected from real users must be carefully anonymised to avoid revealing private information. Moreover, standard smartphone applications typically have limited access to the information about the MBB connection state, thus the context information around this type of measurements is often lacking. Location, type of UE, type of subscription, and RAT are examples of useful information when analysing the results. Our approach uses dedicated measurement

devices, letting us to perform continuous measurements, measure multiple MBB networks simultaneously, and to avoid or filter out any bias in the measurements results due to the hardware and software problems with our nodes or backend.

3.4.3 Measurements using network-side data

MBB performance can also be assessed through network-side monitoring and logging. Network-side logs can give a very detailed insight into the causes of observed behaviour. Network-side logging is typically realised either by tracking a set of UEs on existing CN components, or by installing a network probe within the RAN or the CN.

A main challenge is, of course, that such data is generally only available to operators or their partners. This limitation is one of the main reasons behind the appearance of several testbeds that include the MBB network components as part of their infrastructure. One of the most prominent examples of such testbeds is PhantomNet – a mobile networking testbed developed at the University of Utah [17]. PhantomNet runs in a controlled lab environment and provides access to EPC and RAN components as well as to the set of commercial and Software Defined Radio (SDR)-based UEs. It includes both open source components such as *OpenLTE* and *Open Air interface* as well as proprietary LTE stacks, such as *ip.access* and *OpenEPC*.

A related approach to network-side data is to route user traffic via one or more middleboxes. One of such example is Meddle [79], which proposes to use a combination of middle boxes and Virtual private networks (VPNs) for measuring mobile traffic. Users are envisioned to subscribe to services like ad-blocking and hence allow their traffic to go through Meddle infrastructure. Meddle is an interesting approach, however, it would potentially face a myriad of privacy challenges. It does not accommodate active measurements.

Several measurements used network-side logs to assess various aspect of MBB performance. Gember et al. [35] conducted a large scale measurement study using data from a large US MBB network and from hundreds of controlled experiments with an aim to understand the impact of user activity. The authors showed that network performance can vary significantly depending on user's traffic patterns and location. The authors also compared packet loss on idle and near active devices and found loss rates on idle devices to be 26% higher and likely to be caused by differences between cell sectors. Qian et al. [76] investigated the detailed 3G RRC state machine behaviour and its impact on application performance using TCP

header packet traces from GGSN. The authors found that the RRC state machine may cause considerable performance degradation such as long latencies due to the state promotion overhead. By using TCP-level packet retransmission rates, the authors also found that RTTs and loss rates were correlated, suggesting a connection between time dependent factors and the size of a flow. Shafiq et al. [86] characterised the properties of cellular traffic by analysing traffic records belonging to users in one US state, captured at the CN of one MBB network. The authors showed that the aggregate traffic distribution is highly skewed across different applications and mobile devices, and that only a few applications are popular. Halepovic et al. [39] modelled the time-to-first-byte of HTTP transactions and validated the model using packet traces captured in between SGSN and GGSN in a large US MBB network. The authors found that the time-to-first-byte is a major component in the overall user response time. Shafiq et al. [87] characterised the operational performance of a major US network during two high-profile crowded events in 2012 using voice and data traces collected from the RAN and CN. The authors found that MBB performance degrades during the crowded events when compared to routine days and suggest several mechanisms, such as radio resource allocation tuning and opportunistic connection sharing, to improve the performance. Huang et al. [44] conducted a large-scale study showing the impacts of protocols and applications on LTE network performance. The authors combined controlled experiments with the traces collected from a proxy in the CN of one US LTE network and showed that many TCP connections under-utilise the available bandwidth due to application behaviour and TCP settings. [58] studied the uplink delays in one Austrian UMTS network by conducting controlled set of measurements using a laptop and a USB modem and tracing the packets in the RAN and CN. The authors measured contribution of UE, NodeB and RNC to the uplink delay and derived delay models both for the individual radio and network interfaces as well as for the accumulated delay. Ricciato et al. [80] reported traffic and delay patterns of an Austrian 3G network by using traces from the CN. The authors found that MBB traffic can suffer from large packet bursts generated by scanning activities at high rates in the Internet.

Ability to capture certain signalling messages or user payloads from within the RAN or CN can be indispensable when it comes to precisely measuring inter-arrival times or similar metrics as well as localising MBB performance problems. However, network-side logs can only capture the performance as seen by the network components, but not necessarily as experienced by the user. For example, packet loss due to short coverage gaps can normally only be detected when tracing

the connection locally on the UE. We believe that our end-to-end approach using dedicated measurement infrastructure and long-term measurements can better reveal the MBB performance over time as experienced by the user.

3.4.4 Other approaches to measure MBB networks

A related approach to crowd-sourcing is to piggyback on popular applications for gathering data on various MBB properties. This approach was for example used for geo-locating IP addresses in MBB networks [95]. The authors instrumented a popular iOS application to report the device's local IP address and were able to collect several thousand mobile device locations and IP addresses.

Mobile virtual network operator approach was used by the SciWiNet test-bed [25]. SciWiNet was the National Science Foundation (NSF) project run by Clemson University and provided research infrastructure to data services on Sprint's MBB network infrastructure. SciWiNet maintained a pool of shared smartphones and USB LTE dongles in addition to an Android application that was measuring TCP and UDP throughput, packet loss, and latency.

Several works used infrastructures that were built and deployed for the specific study of MBB performance. One such example is [55] that compared performance of four MBB operators in seven locations in India using several low cost netbooks and USB modems.

3.4.5 Related MBB reliability and performance studies

Perala et al. [74] developed an RRC transition triggering tool and showed that operators do not necessarily configure their RRC state machine in accordance with standard literature. Chen et al. [19] developed a tool called *QoE doctor* to measure mobile QoE with a focus on user-perceived latency, and identify root causes of QoE problems. The authors find that a significant fraction of QoE degradations are caused by network bandwidth throttling and RRC state transition in 3G networks with more than two RRC states. Rosen et al. [81] investigated increased application layer packet delays during RRC state demotions. The authors also discovered increased packet loss during state transitions, but in LTE networks only. The authors conclude that increased delays are mostly due to the issues with state demotion implementations by carriers and therefore suggest that sending packets around state demotions should be avoided. Chen et al. [21] studied RNC-level performance with an aim to understand main driving factors. The authors find

that RTTs can vary widely according to geographical location, coverage of the NodeB and distance between the NodeB and RNC. The authors also observe that in some cases loss is caused by NodeBs during highly loaded periods. Xu et al. [100] combined controlled experiments and crowd-sourced data and showed that downlink traffic can be buffered in the network, causing bursty packet arrivals. In this context, the authors additionally investigated a drop policy enforced by the network and discovered that the drop-tail policy is typically used to drop packets. The authors also found that large downlink buffers typically deployed in MBB networks can cause high latency when throughput is too low to drain the buffers fast. Li et al. [62] studied TCP performance in HSPA+ networks on high-speed rails and showed that the number of handovers is proportional to the increased loss rates for high speeds. Tso et al. [96] measured HSPA performance on the move to be greatly different from static HSPA performance. The authors observed that the final results of handovers are often unpredictable and that UDP packet loss at least doubles during handover periods. Balachandran et al. [12] showed that most HTTP sessions with inter-RAT handovers are abandoned.

The aforementioned studies considered different aspects of MBB reliability and performance. To the best of our knowledge, there has been no comprehensive study of MBB reliability using end-to-end measurements from a dedicated measurement infrastructure. Our work captures failures on the control, data and application levels and correlates them with the connections' metadata. This allows us speculate where in the network failures occur and provide possible root causes with fewer artefacts. One of the main goals of this thesis is to be able to assess reliability in MBB networks as experienced by the user. We think our approach is quite fundamental and unique because we aim to measure the perceived quality and reliability from the end user on a large scale and across multiple operators. We are also not aware of a comprehensive study that characterises reliability in 3G and LTE networks and compares the mobility and stationary scenarios.

3.5 Summary

This chapter has presented a framework for measuring reliability in MBB networks and reviewed the related work. The proposed framework is based on the position that end-to-end measurements can give useful insights about performance and stability problems in the network as a whole. The main argument in the proposed approach is that reliability must be measured at several levels, from the stability

of the network connection to the reliability of the data plane and application layer performance. We believe that this framework gives a good basis for describing the overall reliability of a MBB network.

While the proposed framework was designed for dedicated stationary nodes, it is also applicable for studies based on crowd-sourced data from mobile phones. Such approaches will, nevertheless, often be more limited in the availability of metadata, and in the ability to gather long uninterrupted time series under stable conditions.

Our framework focuses on the set of the most important metrics at each level. There are still many other metrics that can be relevant for understanding reliability. For example, on the connection level, the ability to *establish* a PDP context (EPS bearer) when needed is an important aspect of reliability, which is different than the ability to maintain the connection uninterrupted for a long time. An important topic is also to look at the reliability and performance of various transport layer protocols in MBB networks.

Chapter 4

The Nornet Edge platform for mobile broadband measurements

This chapter describes NNE, a dedicated infrastructure for measurements and experimentation in MBB networks. NNE is the largest infrastructure of its kind, with dedicated measurement nodes distributed in over 100 Norwegian municipalities. In addition to stationary nodes, mobile nodes were introduced by deploying them on regional and inter-city trains in Norway.

In the remainder of this chapter, we highlight the essential features of NNE, present the architecture and describe the main components of the infrastructure.

4.1 Essential features

The main features of NNE are:

Unprecedented scale. In year 2014, NNE consisted of more than 400 measurement nodes. The large number of nodes made it possible to give a representative view of the characteristics of an entire network.

Nationwide geographical coverage. The measurement nodes are widely distributed across Norway (including one node at Svalbard). The challenging Norwegian topography makes it possible to collect measurements under diverse conditions, from major cities to remote islands. At the same time, there is a dense deployment of nodes in a few main cities, giving a more detailed view of network conditions in urban areas.

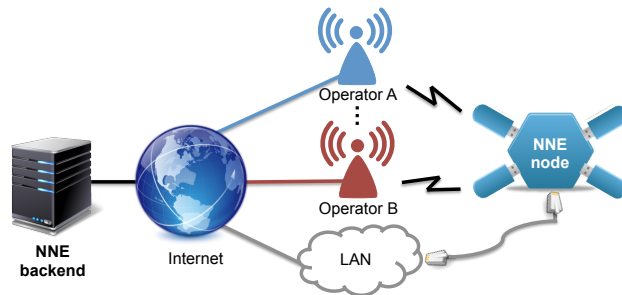


Figure 4.1: NNE overview.

Fully programmable measurement nodes. The measurement nodes are embedded computers running a standard Linux distribution. They are flexible and powerful enough to run most measurement tasks, including video experiments.

Multi-homed measurement nodes. All NNE nodes were connected to at least two MBB providers, and often also to fixed or wireless networks. This made NNE particularly well suited for experimentation with methods that exploit multiple communication links.

Rich context information. In addition to information about network, time and location for a measurement experiment, NNE nodes have built-in support for collecting information such as CID, signal strength, connection mode etc.

Advanced system for experiment management. NNE makes it easy to deploy experiments on all or a selected subset of the nodes. NNE also supports transferring and storing measurement results in a central sever with minimal effort from the experiment developer.

4.2 System overview

The main goal of NNE is to measure end-to-end performance as seen by the end user. This performance is influenced by properties of the operating system, the MBB modems, the RAN, the underlying transport network, and the mobile CN. NNE allows direct comparison of different MBB networks, by using the same hardware and operating system for all providers.

Figure 4.1 shows a schematic overview of the NNE infrastructure. It consists of two main components: a large set of NNE nodes, and a central backend system. The NNE nodes run a standard Debian Linux distribution and a patched Linux 3.x kernel, giving great flexibility in the types of measurements that can be supported. The NNE nodes are described in detail in Sec. 4.3. The backend system consists of a number of servers for monitoring and controlling the nodes, deploying and managing measurement experiments, and processing, storing and visualising measurement data. The backend contains an Secure Shell (SSH) proxy server that enables remote login to NNE nodes. The backend system is described in detail in Sec. 4.4.

Measurement data often becomes much more valuable when combined with additional information about the context in which they were collected. For example, it is often interesting to know the CID of the base station that the node is connected to, the signal strength, the RAT of the connection (e.g., 2G, 3G or 4G), the state of the Radio Resource Controller (RRC) on the node etc. NNE nodes contain a built-in tool that collects this information from the MBB modems, and makes it available for measurement applications in a publish-subscribe fashion.

4.2.1 Deployment of nodes

NNE nodes are distributed to reflect the population density in Norway, with some bias towards urban areas. Nodes are placed indoors in small or large population centres, with a higher density of nodes in larger cities.

NNE nodes were primarily deployed in collaboration with Norwegian communes and the Ministry for Local Government and Regional affairs. Their interest in NNE was related to electronic voting. Voter registration is done electronically, and all voting centrals must have a working Internet connection. NNE nodes were used to monitor the existing Internet connection in these voting locations, and to provide a robust backup connection in case the primary connection fails. Hence, from the start until year 2014, most NNE nodes were placed in voting locations in Norwegian communes. Voting locations were often schools, city halls or other public convention centres. Five large cities had NNE nodes placed in all voting locations. Back in 2014, 289 of our 443 nodes were placed in these 5 cities, with 136 of these in the largest city Oslo. Overall, more than half NNE nodes were deployed in three largest cities, where 26.7%¹ of the country population lived. 127

¹<http://www.ssb.no/en/befteft/>

of the remaining nodes were spread across 85 other communes, with 1 - 6 nodes per commune.

Figure 4.2 shows the placement of NNE nodes in Norway as of June 2014, classified according to the number of MBB networks the NNE node was connected to.

4.2.2 Multi-homing

A particular focus in NNE is to support multi-homing, enabling multiple connections to be used in parallel. Each NNE node was connected to at least 2 and up to 5 MBB providers, using standard subscriptions. In addition, the nodes had a fixed Internet connection when this was possible. Multi-homing opens up many opportunities for novel types of experiments. At the infrastructure layer, NNE can be used to discover correlations in the performance characteristics of the different MBB networks. Such correlations are important for applications that rely on multi-homing for increased robustness or performance. At the transport layer, NNE is well suited for experimentation with novel protocols such as Stream Control Transmission Protocol (SCTP) [93] and MultiPath TCP (MPTCP) [33], which allow concurrent transmission over multiple paths.

4.3 The Nornet Edge node

This section discusses the requirements that led to the selection of the current hardware platform, and gives a description of the hardware and software functions that the node offers to support experiments.

NNE nodes are small embedded computers running a standard Linux distribution, and the software and tools that run on the nodes are all based on standard Linux libraries. Any computer that can run Linux can therefore in theory be used as an NNE node. We have, however, chosen to stick to a single hardware configuration for our nodes. This makes management and administration of the nodes easier. It is not unlikely that the situation with a single hardware platform will change in the future, when new generations of NNE nodes will be deployed.

Several requirements were taken into consideration when selecting the hardware and software configuration for NNE nodes, including price, performance, availability and user community. Based on these requirements, several differ-

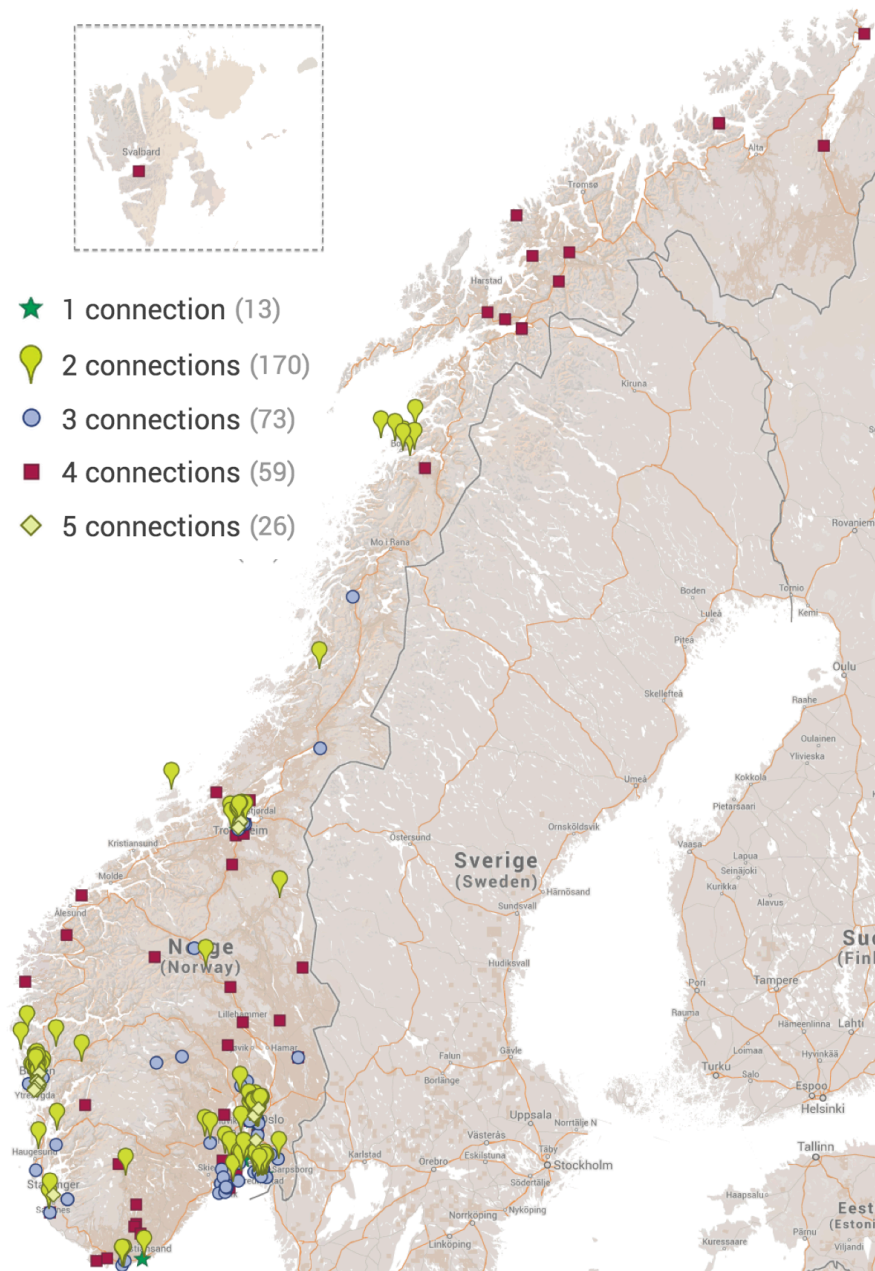


Figure 4.2: Placement of NNE nodes in Norway as of June 2014.

ent single-board computers were tested, including the BeagleBoard², the Beagle-Bone³, the PandaBoard⁴ and the Raspberry Pi⁵. After extensive testing, it was found that none of these offer the stability and performance required. The USB controller of most of the boards was unfortunately not able to cope with multiple USB modems connected to it. In addition, none of the tested boards had a sufficient number of built-in USB ports. This required the use of an external USB hub, with its own power supply in order to give enough power to the connected USB modems.

Based on this experience, it was decided to develop a new single-board computer, specifically made to be a hardware platform for NNE. This may sound like a radical decision, but it was deemed necessary in order to have a sufficiently stable node. Hence, the *UFO-board* (named after its shape) was designed and produced in collaboration with Dynsense⁶, a Hong Kong based hardware development company.

The main drawback of designing a board specially for this project is the lack of an existing user community. Note, however, that most of the components on the UFO-board (CPU, USB controller, Ethernet port, etc.) are off-the-shelf components that are supported by the mainline Linux kernel. Only a small patch is required to make the UFO-board run. The advantage of having a custom designed board is the ability to select components that fit our needs. This includes a CPU that is powerful enough to support multiple parallel experiments on different network interfaces, enough USB ports to support a large number of MBB modems and other peripherals, and a USB controller and power supply that can handle many connected devices. So far, the UFO-board has proven to be well suited for its purpose, and the stability has been far better than for any other tested board.

4.3.1 Technical specifications

Figure 4.3 shows a picture of an NNE node. The exact configuration of the nodes varied slightly during the different phases of deployment. The main components of it are the following:

UFO-board single-board computer. The UFO-board has a Samsung S5PV210

²<http://beagleboard.org/>

³<http://beagleboard.org/bone>

⁴<http://pandaboard.org/>

⁵<http://www.raspberrypi.org/>

⁶<http://dysense.com/>

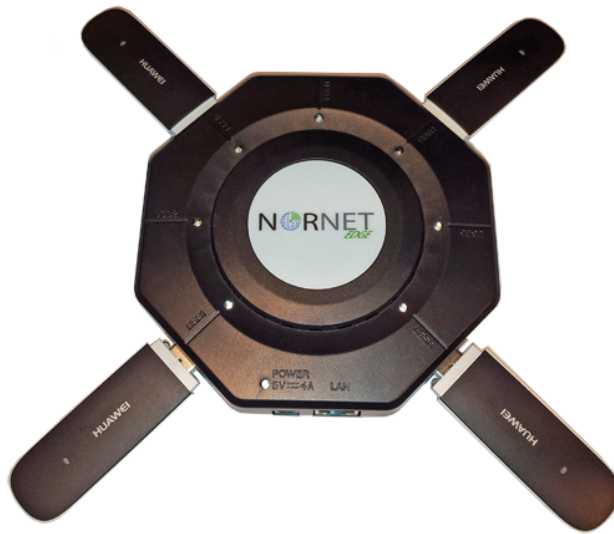


Figure 4.3: NNE node with four USB modems.

Cortex A8 1GHz processor. It has 512 MB RAM and 512 MB NAND Flash memory, and a 16GB SD card for storage. 7 USB ports and 1 Fast Ethernet port is embedded on the device.

Linux operating system. The UFO-board runs a standard Linux distribution, initially with a 3.0.8 kernel, later upgraded to version 3.8.8 and then to version 3.10.39. The root file system includes standard Debian Wheezy components as well as the software described below. The image including the root file system is written to the SD card during the assembly process. During this process, each UFO-board is also given a unique node ID.

1 - 4 3G (UMTS) or 4G (LTE) modems. Huawei E353-u2 (3G) and Huawei E392-u12 (LTE) USB modems were used for all UMTS and LTE operators, respectively. The former modem supports GSM technologies up to HSPA++ ("3.5G"), whereas the latter supports up to LTE Category 3. A main motivation for selecting these models was that they were sold to end-users by some of the operators that the nodes connect to, which ensures that our measurements are comparable to a typical user experience. Another advantage of these modems is that they report about events such as change of RAT, CID, signal strength, and RRC state. In year 2015, all Huawei E353-u2 modems were replaced with Huawei E392-u12. The same modem model is always

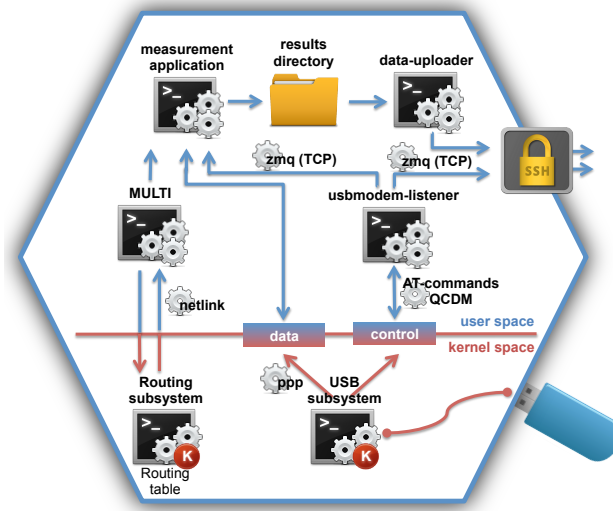


Figure 4.4: The most important functions on the NNE node.

used for all networks on the same node, to avoid differences caused by different hardware.

1 CDMA (1x Evolution-Data Optimized (EV-DO)) modem. This modem was used to connect to Ice, which operates a MBB (data only) network in the 450 MHz frequency band. Two different modem models were used, supporting the Rev. A and Rev. B versions of the 1x EV-DO standard respectively.

Wi-Fi module. Some NNE nodes were equipped with a TP-LINK TL-WN822N USB Wi-Fi module. This module served two purposes. First, it allowed the NNE node to connect to an available Wi-Fi network. Second, it was used to turn the NNE node into a Wi-Fi access point, providing Internet access through the MBB connections.

4.3.2 Management and support functions

NNE nodes contain a set of tools and services that are useful for many measurement experiments. These support functions give experiment developers easy access to information about the state of the MBB connections, and perform standard operations like uploading measurement results. Figure 4.4 shows a schematic overview of the most central components running on an NNE node.

SSH connection

NNE nodes can be reached remotely via SSH. Since most nodes do not have a fixed IP address, the SSH connection is initiated from the node. This is done using *autossh*, which automatically (re-)establishes a reverse SSH connection to our server on a port deduced from the unique node ID.

The SSH connection will normally be established over the default route on the node (i.e., the route with the lowest metric). The default route is set individually on each node, based on the observed quality of the different connections (the fixed Local Area Network (LAN) connection is always preferred if it is present). *autossh* will retry a limited number of times to re-establish the connection over the default route if it goes down. If not successful, it will give up and exit. The *autossh* process will then be restarted and will establish the connection over the next available connection with the lowest metric.

usbmodem-listener

The *usbmodem-listener* is a daemon written in Python that incorporates three main functions: modem management, cellular connection management and metadata collection.

The MBB modems used in NNE nodes are mostly USB based. These modems expose multiple serial devices to the operating system. Before the modem can be used for communication, a series of AT commands must be sent on the correct serial port in order to identify the SIM card and establish the connection to the network. *usbmodem-listener* listens to *udev* events from the Linux kernel, and performs the necessary configuration as soon as a new USB modem is plugged in. It spawns the Point-to-Point Protocol daemon (PPPD) process that creates the Point-to-Point Protocol (PPP) connection to the modem. *usbmodem-listener* keeps monitoring the state of the forked *pppd* process and restarts it when it exists or becomes stale. It also makes sure that all occupied resources are freed when the modem is disconnected.

The performance of a MBB connection is very often correlated with its RSSI, RRC state and other state variables. To better explain the observed behaviour, it is therefore important to collect state changes along with measurement results. To this end, *usbmodem-listener* collects metadata information about each active network connection and publishes it to subscribers. NNE nodes use Huawei modems with a Qualcomm chipset. We extract a large number of status variables from the

respective serial ports of the modems using AT commands and special tools developed for Qualcomm-based modems [63, 78]. The state attributes that are the most relevant to our measurements are connection RAT (GSM/GPRS, WCDMA, LTE), radio signal information (RSSI) and signal-to-noise ratio (E_c/I_o), RRC state and camping network operator. In addition, we also record when a connection comes up or disappears, i.e., when the PDP context (EPS bearer) is established or lost. Since access to the serial ports is limited to a single process at a time, the *usbmodem-listener* acts as a publisher of the status variables it receives from the modems, and makes it available to local or remote subscribers through a socket interface.

MULTI

NNE nodes use MULTI [30] to manage the network connections. MULTI is a command line network manager for Linux with support for multi-homing. The currently available network managers for Linux do not properly configure routing tables with multiple active interfaces. Without special consideration, multiple interfaces cannot be used in parallel.

MULTI automatically detects when interfaces are connected/removed, and then obtains an IP address (if needed) and configures the routing tables correctly. After routes have been added or removed, MULTI broadcasts a notification to a MULTI Netlink group. Netlink is a socket-based, Inter-process communication (IPC) mechanism for Linux. Applications can subscribe to these messages and adapt based on changes in network state. For example, a measurement application can be configured to automatically start measuring on interfaces when they become available.

data-uploader

NNE offers functionality to ease the process of uploading measurement data to a server. Measurement applications write their output in a predefined format to a special output directory on the local file system. The *data-uploader* will periodically scan this directory and transfer log files to a corresponding directory at the server-side data collector described in Sec. 4.4. This approach simplifies near real-time visualisation of measurement results, without having to write special upload and encoding/decoding procedures and importing data into a database in each measurement experiment.

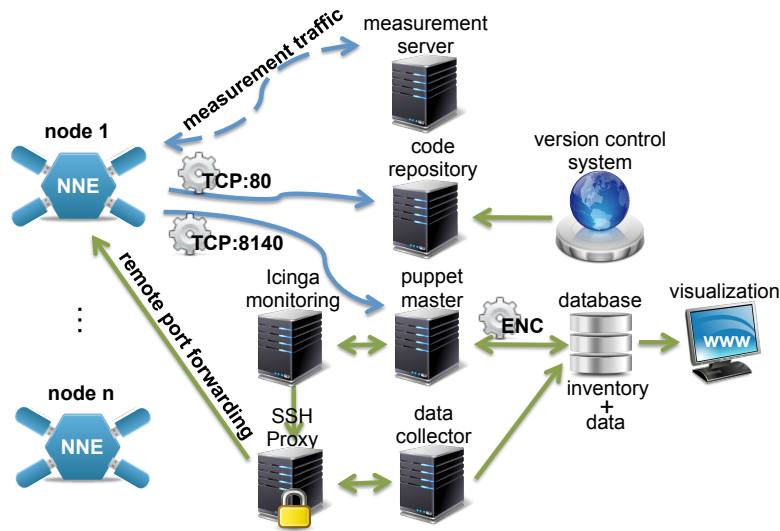


Figure 4.5: NNE backend.

4.4 Backend system

The NNE backend system consists of a number of servers that collect, store, process and visualise measurement data, and perform various monitoring and control tasks for the NNE nodes. The NNE backend is based on a VMWare cluster, with an attached redundant array of independent (RAID) shelf for storage. Figure 4.5 shows an overview of the NNE backend with its various functions.

4.4.1 Node control and management

SSH proxy

Each NNE node maintains an SSH connection to a server acting as an SSH proxy, which makes it possible to access the nodes remotely. The SSH proxy accommodates both local and remote port forwarding to the nodes. Local port forwarding is used on a node if a restrictive firewall policy is enforced by a LAN connection, so that the node cannot reach the backend servers directly. Remote port forwarding is used so in order to make the SSH server and the updates produced by the *usbmodem-listener* on the NNE nodes accessible remotely. Remote access to the SSH proxy itself and the forwarded ports is limited only to the NNE backend servers that need to reach the nodes.

Puppet for node management

Puppet⁷ is used for managing and updating NNE nodes. The main purpose of Puppet is to push changes and updates in measurement or management software to NNE nodes automatically. The Puppet infrastructure consists of Puppet agents running on each node, that connect to a central Puppet master. NNE nodes are preconfigured to start the Puppet agent on boot, which then connects to the Puppet master. The master compiles the configuration for the node based on manifest files, and sends it to the agent. The agent in turn checks the received configuration against the current system state, and takes the necessary action by creating or updating resources like files, Debian packages or configurations. It then reports the results back to the master and repeats the whole procedure every hour.

Icinga for monitoring

Icinga⁸, a branch of the well-known Nagios monitoring system, is used to perform periodic availability and service checking of NNE nodes. Each check is done by connecting to the node over SSH via the SSH proxy server, and reporting the status of parameters like disk space, processor load etc.

Build server and package repository

NNE has a simple process for deploying and updating measurement applications and other software. When a new version of an application is ready, it is pulled from the version control system and placed on the package repository server. Most measurement applications and other software that is installed on NNE nodes are packaged as Debian packages for ease of management, but they can also be distributed in other forms of software packaging systems. Once a new package (or version) is available in the code repository, the Puppet master will instruct the relevant nodes to download and install the update.

Measurement server

Some measurement experiments require active participation from the server side. The NNE backend contains a measurement server that can be used as an endpoint for these experiments. The measurement servers are well provisioned in terms of

⁷<https://puppetlabs.com/>

⁸<https://www.icinga.org>

bandwidth and processing power, to make sure they are not a performance limiting factor. Experiments can also be run against any other server on the Internet.

Data collector

Measurement data collected at the nodes is transferred to the backend by the *data-uploader* described in the previous section. At the backend, the files are received and processed by the data collector, which inserts data into a relational database. The data collector can insert the data directly, or perform some processing before the data is inserted, like calculating averages or extracting selected values.

4.4.2 Database

The database is a central component in the NNE backend, that serves two main purposes. First, it keeps records for the NNE inventory, consisting nodes and their associated components like modems, subscriptions and other hardware. For each node, there is also information about location, address, contact person, administrative messages etc. Second, the database is used to store data from the measurement experiments. By using a single database to store both inventory information and measurement result, it becomes easy to extract measurement results from a particular node, and to compare results across nodes.

Measurement data is usually stored in a generic database table, which can be used for any type of measurement. The table contains a timestamp, pointer to the measurement instance (e.g., the node, network and type of measurement), as well as the output from the measurement in an Extensible Markup Language (XML) format. The format of the measurement data is specified in a separate table, and can for example be an RTT value for a single data packet, the setup time for a TCP connection, or the result from a throughput test. The use of a generic measurement data table makes it easier to deploy new types of measurement experiments. In addition to the raw measurement data, measurement results are post-processed to calculate aggregates and also to filter out time periods from problematic connections, NNE maintenance windows or when NNE experienced problems at the server-side due to hardware problems or problems with their network provider.

With a large number of nodes running continuous measurements on up to 5 MBB interfaces in parallel, the database must be organised in a scalable way. After considering several options for scalable data storage, such as Apache HBase, we finally decided to use MySQL and stay with a relational database model. This

decision was made because we have a fixed scale deployment, where the number of nodes, networks and approximate number of daily database records is known in advance and will not grow uncontrolled. To improve scalability, the dataset is split into smaller chunks using the MySQL partitioning feature. Daily partitions are created for the tables that contain large amounts of measurement results. Partitioning also facilitates fetching of a range of data fast, since it only involves scanning partitions that fall into a given interval of dates, so for example in order to select data for past 15 minutes, a scan of only one partition is made.

4.4.3 Visualisation

The status of NNE nodes and results from selected measurement experiments are visualised on a website associated with the Nornet project. The website shows a real-time view of the status of all NNE nodes, including the status of each MBB connection. The status of the different connections are visualised as colour codes on a map, and status can be shown for individual operators or all operators combined. A screenshot from the website is shown in Fig. 4.6. The status information is collected by the *usbmodem-listener* as described in the previous section, and status updates are published to a status collector running in the backend in real time, so that the website is updated within few seconds.

By clicking on the nodes in the map, more detailed information about the current status of each connection is shown, as well as graphical representations of various measurement results. Results from some measurement applications are also shown in near real-time, with a delay of approximately 1 minute. The liveness of such data is determined by how frequently the application writes its output to the results directory on the node, and how frequently the *data-uploader* scans this folder and transfers the data to the server. In addition to (near-)realtime data, the website also shows historical data on how different metrics have evolved over time at the different nodes.

The visualisation website serves several purposes. First, it is an efficient way for researchers to navigate in the collected data. Second, the website serves as a point of entry for media and the general public who are interested in assessing and comparing the performance of the different MBB operators, and in seeing data on how various failures and incidents in the network affect performance. Last, the website serves as an easy way for our collaborators (operators and node hosts) to get a real-time overview of the status of their network and measurement nodes. The operators use a tailored version of this website to get a view of the

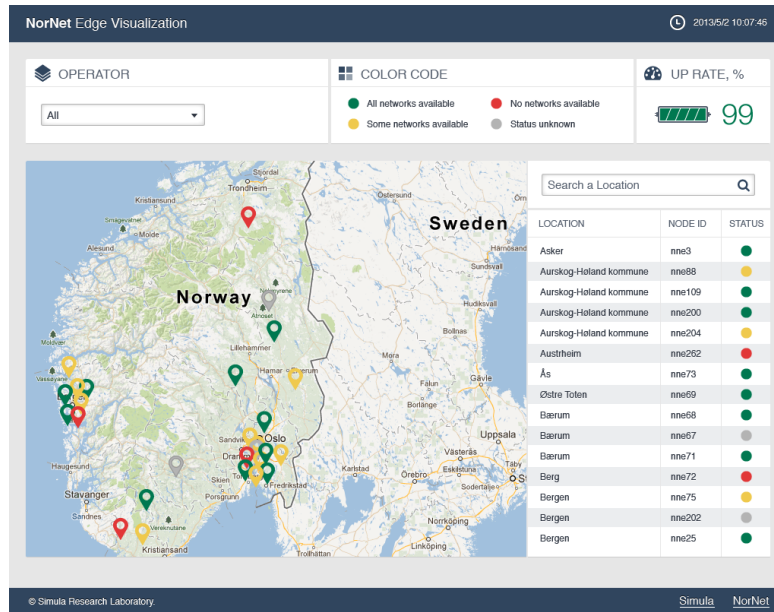


Figure 4.6: Screenshot from the NNE visualisation website.

customer-experienced performance of their networks, to complement their own network monitoring.

4.5 Nodes on trains

The first mobile NNE nodes were first deployed on regional and inter-city NSB trains in the end of year 2014. NSB is a government-owned railway company that operates most passenger train services in Norway. Like other NNE nodes, train nodes connected to up to four UMTS and LTE operators via Huawei E392-u12 modems that support up to LTE Category 3, and one CDMA 1xEV-DO operator. We use the Global Positioning System (GPS) location and speed data from the train's fleet management system to identify the location of NNE measurement nodes and trains speed during the measurements. The GPS locations are updated every 10 to 15 seconds in the fleet management system.

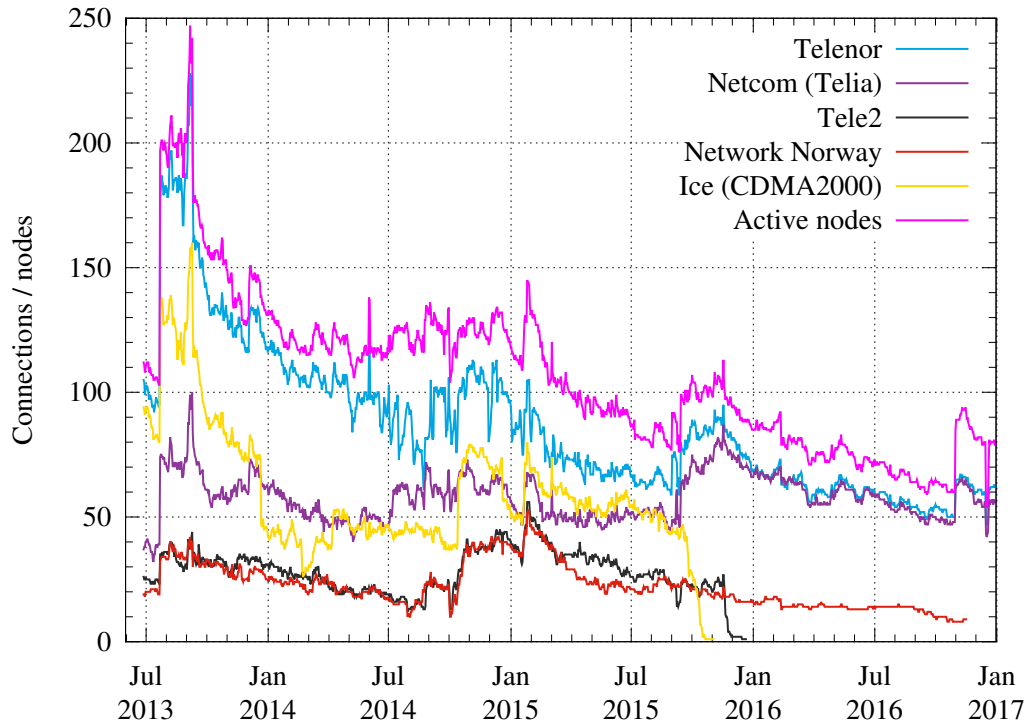


Figure 4.7: The number of active NNE nodes and MBB connections over time. Note that the counters for Tele2 and Ice connections drop to zero in 2015 because of the shut down of the respective networks.

4.6 Challenges with running a large-scale measurement infrastructure

4.6.1 Maintaining high number of active nodes and MBB subscriptions

In the beginning of 2012, all existing MBB operators were invited to take part in the establishment of NNE, and three of them entered into a collaboration on building NNE. These three operators offered free MBB subscriptions to NNE nodes. For the remaining two operators, we bought subscriptions. The operators' main motivation for taking part in NNE was access to real-time data on the current status of all MBB networks as seen by the end users. This information was seen as useful by the operators to complement their own network monitoring.

The collaboration with the Norwegian municipalities has proven to be a good way to manage the physical distribution of NNE nodes. It gave a good geographical spread of the nodes, which was representative for the population in Norway. Local IT support at the communes was responsible for installing the nodes, and also helped with on-site management such as physically rebooting a node or replacing faulty equipment when needed. Most management of NNE nodes was done remotely (see Section 4.4), but occasionally there was a need for physical intervention, e.g., when contact with a node is lost.

Figure 4.7 shows the number of active NNE nodes and MBB connections over time. The varying number of nodes and connections has several underlying reasons. First, the number of nodes has reached its peak of around 250 in September 2013 - the time the Norwegian Parliament election was held. After the election, a significant fraction of municipalities decided to discontinue hosting NNE nodes. This triggered an internal process of finding new hosting partners, but this appeared to be challenging and slow with respect to the number of new hosts acquired. Second, a large number of nodes were lost over time. In some cases this was because hardware or software problems, and normally we contacted the host and suggested replacing the node. In some other cases, however, we lost both the connection to the node and the contact with the host.

4.6.2 Typical failures and preventive actions

From the first phase of NNE deployment, the stability of the NNE node was of top priority. Since the nodes are head-less (no display is attached to them), all troubleshooting has to be done remotely. This implies that in order for us to be able to troubleshoot a node, it needs to have at least one working MBB or LAN connection and an established SSH tunnel towards our SSH proxy. Our experience showed that even at the locations with good MBB coverage, the MBB connections tend to fail from time to time, and in some cases failures can be long-lasting if no recovery actions are taken. It was therefore critical to continuously update the software running on the NNE nodes that they would automatically recover from newly discovered failure patterns. In addition to the automated recovery, we have a watchdog process running on all nodes, which attempts a soft reboot of a node if it is unable to establish the SSH tunnel towards our SSH proxy in one hour. In the following, we share some of the failure situations we have experienced and the preventive actions taken.

Stuck USB modems. After the initial deployment of NNE nodes, we noticed that

some of our connections went down and were not able to recover automatically. This happened despite the fact that our connection manager, *usbmodem-listener*, was periodically checking the state of MBB connection and re-attaching to the network when necessary. In case re-attachment fails, *usbmodem-listener* attempts to reset the modem. After a set of manual inspections we realised that in some cases the USB modem becomes unresponsive, hence all standard commands to reset the modem fail.

We contacted Huawei, the vendor of our modems, and provided all necessary information about the situation we encountered with. Huawei provided a new firmware that we flashed to all modems not yet shipped out to the hosts. However, there already was a large number of nodes with the modems running older firmware attached. The solution was found after we realised that whenever the modem gets stuck, it switches into a so called download mode, which has a distinct USB product identifier. From that moment on, the modem only accepts binary commands from Qualcomm's download mode specification. Fortunately, one of the commands defined is to reset the modem. We added the execution of this command to a *usbmodem-listener* monitoring procedure upon detection of the USB mode change.

The cases like this shows that USB MBB modems are not always stable and the recovery actions can be non trivial. It is therefore difficult to foresee such situations in advance. One of the solutions could be the ability to power-cycle individual USB ports (hence the connected modems) or the whole node. The NNE nodes described in this chapter did not have the power-cycling feature, but it is present on the new generation NNE nodes.

Stale MBB connections. The aforementioned connection and modem monitoring procedure considers the connection as up and running if the following conditions are met:

1. the modem is responding to AT commands;
2. the modem is attached to the network;
3. there is a corresponding PPPd process forked;
4. there is a corresponding PPP network interface with an IP address assigned.

Only if one or more conditions are not met, the procedure starts its set of actions to recover the connection. In 2015 we discovered that individual Telenor connections were seen as up and running, but the UDP pings to our echo server resulted

in 100% packet loss. Detailed analysis of such connections revealed that after a few days of normal operations, IP traffic started to be treated differently. More specifically, UDP traffic was blocked, but some TCP packets were still going through.

Additional inspection showed that there was a TCP proxy in between, terminating and proxying all TCP connections. The treatment of the traffic became normal again upon the re-establishment of the PDP context (EPS bearer). We can only speculate what led to such treatment of the traffic. One explanation could be the fact that one of the CN nodes lost the state of the PDP context (EPS bearer) and from that moment on packets were redirected to a middlebox that is normally used for other purposes such as deep packet inspection.

The described situation may lead to an unreachable remote measurement node if the affected MBB connection is the only available Internet connection on that node. From the infrastructure and node management point of view, stale connections should be recovered automatically whenever and as soon as possible. On the other hand, forceful connection recovery may create certain bias with respect to capturing the MBB network misbehaviour. It is therefore important to have at least two working network connections on a single node as well as to not start connection reset too fast, so that the misbehaviour is detected before the connection is re-established.

4.7 Summary

This chapter has introduced Nornet Edge, a dedicated infrastructure for measurements and experimentation in MBB networks. NNE consists of a large number of measurement nodes, places in voting locations all over Norway, and a set of servers constituting the backend system. The NNE nodes consist of a special-made single-board computer running a standard Linux distribution, with attached MBB modems from up to 5 different operators. The backend system contains functionality for managing nodes, collecting and storing measurement data, and visualising measurement results.

NNE can be used for a wide range of MBB experiments that target the performance of the MBB network itself, or the performance of protocols and applications running in the network. NNE is particularly well suited for experiments that require long-running or periodic measurements from a large number of identical and controllable measurement nodes, and a large geographical footprint. One of the

unique features of NNE is the possibility to run experiments on multiple MBB networks in parallel and on identical hardware, which allows direct comparison of performance metrics across operators. NNE also automatically collects status information from the modem, so that measurement data can be enriched with information about the cell ID, connection RAT, signal strength etc.

MBB is designed to make it easy to deploy new measurement applications and start collecting data. The infrastructure was partially made available to other researchers and was preserved to be a valuable resource for the network measurements community.

Part II

A large-scale study of reliability of Norwegian MBB networks

Chapter 5

Measured MBB networks and measurement datasets

5.1 Measured MBB networks

The situation in the Norwegian MBB market has changed over the period of this thesis. Until 2015, there were five MBB operators in Norway. Four of these (Telenor, Netcom, Tele2 and Network Norway) were GSM/UMTS/LTE networks, while the last (Ice) was a CDMA2000 network. As shown in Figure 5.1, only three networks, Telenor, Netcom, and Ice, maintained their own nation-wide RAN. Ice ran a data-only CDMA2000 network in the 450 MHz frequency band. This network was qualitatively different from the GSM-based networks, since it used a different technology, operated on a much lower frequency, and used different end-user equipment.

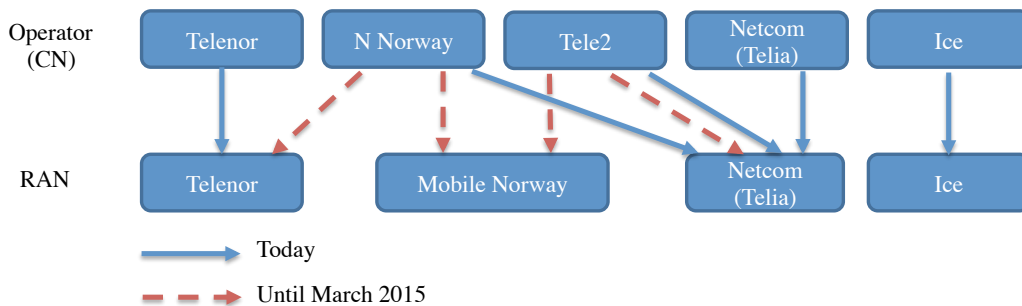


Figure 5.1: The operators and RANs measured in this study.

In 2011, Swedish Tele2 acquired Mobile Norway, the third Norwegian GSM/UMTS operator. Before the acquisition, Network Norway owned 50% of the Mobile Norway and the other 50% of the shares belonged to Tele2. The two owners of Mobile Norway were collaboratively expanding the RAN shared between Tele2 and Network Norway customers, which then did not yet have nation-wide coverage. When outside of their home network, Tele2 customers camped on Netcom's RAN, while Network Norway customers camped on Telenor's RAN. This complex relation between the operators and RANs was an advantage for our measurement study. By looking at correlations between connections on the same RAN but in different operators (or vice versa), we could often determine whether an observed behaviour was caused by the RAN or the CN.

In 2013, Tele2 and Network Norway were merged. In 2014, both Telenor and Netcom started offering commercial LTE services to customers. In 2015, Tele2 and Network Norway were partly acquired by TeliaSonera, an owner of Netcom (which in 2016 has changed its name to Telia), and partly by Ice. In 2015, Ice has started building their LTE network for data-only services in the 450 MHz frequency band and shutting down their old CDMA2000 network.

Regardless of the year of deployment, all NNE nodes were connected to at least two and up to five different MBB networks.

5.2 The effect of mobility on MBB reliability and performance

The studies presented in this part look at the MBB reliability and performance using two different scenarios: when MBB connections are stationary, and when under mobility. The reasons for such a split are both historical and practical. The initial NNE deployment consisted of only stationary nodes, which produced the data for the first studies presented in this thesis. The mobile nodes were only introduced in the later phases of the project. The data collected from mobile nodes not only showed clear differences in reliability of stationary and moving connections, but also revealed a set of challenges, which severely impact connection stability when under mobility. Our study on the the effect of mobility presented in Section 7.1 established the very different loss rates in a mobile vs stationary scenario. Figure 5.2 shows the overall loss rate for the two major Norwegian MBB networks, Telenor and Netcom, when the nodes were stationary and moving, and when the time series are aggregated into 5-minute bins. It is clear from the figure that loss

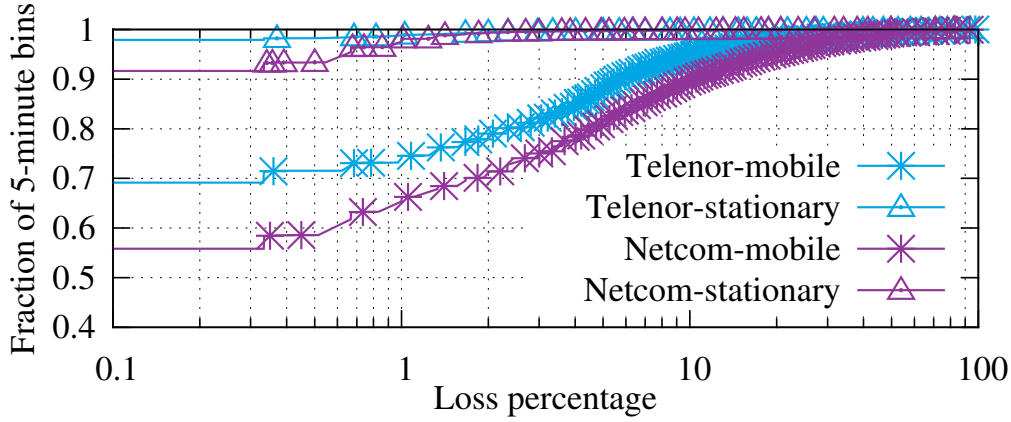


Figure 5.2: Overall loss rate and the effect of mobility. Much higher loss rates observed when mobile in both networks.

is much higher when mobile than when stationary. When the nodes are stationary, only 2 % (Telenor) to 12 % (Netcom) of 5-minute bins involve packet loss. In the mobile case we observe loss in 30% (Telenor) to 50% (Netcom) of bins, and 5% (Telenor) to 10 % (Netcom) of bins have a loss rate above 10%.

The large difference in loss rates between mobile and stationary nodes observed in the later phases of the research presented in this thesis largely motivated to treat and analyse the two scenarios separately based on the MBB connection mobility.

5.3 Measurement datasets

The studies in this part are based on the dataset collected from the stationary and mobile NNE nodes. The stationary nodes were geographically distributed in rural and urban areas of Norway. The mobile nodes were placed on board regional and inter-city trains. Both stationary and mobile nodes were connected to two or more Norwegian MBB networks. Figure 5.3 shows the different datasets used in this thesis and pointers to the respective studies where individual datasets were used. The datasets span three years and consists of the measurement results as well as connection metadata and inventory data.

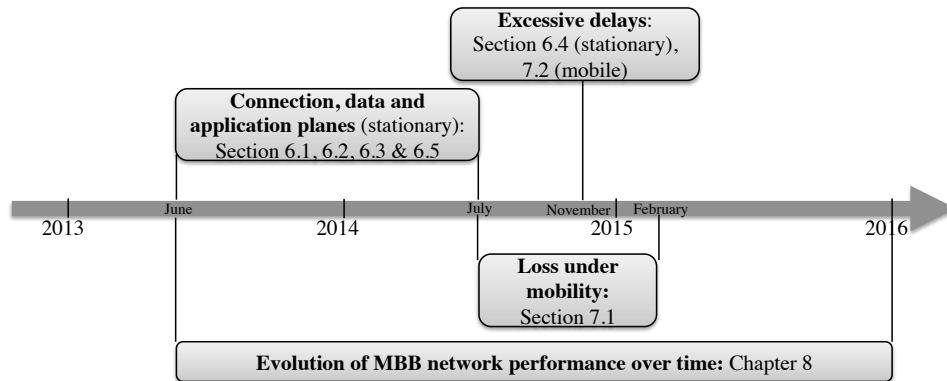


Figure 5.3: Measurement datasets used in the study.

5.3.1 Stationary scenario

The studies in Chapter 6 are based on the dataset collected from the stationary NNE nodes connected to up to five MBB networks in 2013 and 2014. Connection (Section 6.1), application (Section 6.5) and parts of the data plane (Sections 6.2 and 6.3) reliability and performance studies are based on the dataset collected from June 2013 until July 2014.

This dataset consists of more than 13 billion measurement points, gathered from 938 distinct connections at 341 distinct nodes. 327 of these were Telenor connections, 142 were Netcom, 75 were Tele2, 66 were Network Norway, and 328 were Ice¹. The number of simultaneously active measurement nodes has varied in the range between 108 and 253 through the measurement period.

The study on excessive delays (Section 6.4) is based on the data collected from over 200 stationary Telenor and Netcom connections for a period of one month, November 2014.

5.3.2 Mobile scenario

The study of classifying loss under mobility (Section 7.1) is based on data collected by the six measurement nodes placed on trains in Norway from July 2014 until February 2015.

The study on excessive delays (Section 7.2) is supplemented with the mobility scenario that is based on the data collected from four NNE mobile nodes for a

¹The varying number of connections per operator was caused by practical and economical constraints.

period of one month, November 2014, using, among others, Telenor and Netcom MBB connections.

Chapter 6

Reliability and performance: stationary scenario

6.1 Connection reliability

Data can only be sent over a MBB connection when there is an established PDP context (EPS bearer) in the CN. To establish a PDP context (EPS bearer), the UE signals its presence to the respective signalling gateway (SGSN in UMTS networks, MME and SGW in LTE networks), which then establishes the PDP context (EPS bearer) and returns a data session with an allocated IP address. This data session is essentially a tunnel connecting the UE to the Internet through intermediate gateways (GGSN in UMTS networks, PGW in LTE networks). The PDP context (EPS bearer) can be broken either by problems in the RAN (e.g., poor signal quality), or in the CN (e.g., failures or capacity problems in the SGSN). Failures can also be caused by the complex interaction between the OS running on the measurement node, the node's USB subsystem, and the MBB USB modem itself. We conjecture, however, that if the majority of failures are caused by such artefacts, the differences between operators would be minor and hard to spot.

In this section, we measure the frequency of PDP context losses, the time it takes before the PDP context is successfully restored, and the resulting downtime when no PDP context is available. We further correlate with signal quality and RAT to gain an insight into what may have triggered the failure.

The discussion in this section is limited to the UMTS networks, since the study was performed in the first half of 2014, when the NNE nodes were not yet equipped with LTE modems. We also exclude the CDMA2000 network, since the

CDMA2000 used did not provide the necessary logs about the connection state.

6.1.1 Measuring connection failures

An NNE node continuously monitors the status of the PDP context for all UMTS connections, and tries to re-establish it as soon as it breaks. If it fails in doing that, the node keeps retrying until it eventually succeeds; we log all these attempts. There is no hold time between these consecutive reconnection attempts, so a new attempt is immediately initiated after the failure of the preceding attempt. A failure will therefore trigger a varying number of reconnection attempts depending on its duration (each attempt takes tens of milliseconds).

In some cases, the node manages to re-establish the PDP context for a short period, before it breaks again. To build a time series of failure events, we group consecutive reconnection attempts spaced by less than M minutes into the same event. In other words, a connection must keep its PDP context for at least M minutes before the reconnection was deemed successful and the failure event ends. Setting M to a high value underestimates the connection stability, while a low value will report a flapping connection. We experiment with different values for M in the range from 1 to 5 minutes. We detect a total of 154772 failures when setting M to 1 minute. Varying M from 1 minute to 3 minutes has a negligible impact on the number of detected failures. This number only drops by 0.019% when we set M to 3 minutes. Based on this, we set M to 3 minutes when identifying PDP context failures. We believe that this grouping captures well what the user perceives as a usable connection, since a connection is not worth much if it flaps at a high frequency. The result of this grouping is a sequence of connection failure events of varying duration for each connection.

We impose two conditions to avoid overestimating the number and duration of connection failures by including measurement artefacts. First, we discard all failure events that were rectified either by rebooting the node or actively resetting the USB modem, since these may be caused by effects in the measurement platform. Second, to compensate for absent log files and failures that are not rectified by the end of our study period¹, we consider only failures for which we have logged both starting and ending points with no gaps in between.

¹In some cases, measurement nodes were lost for varying periods of time which resulted in gaps in the logged data.

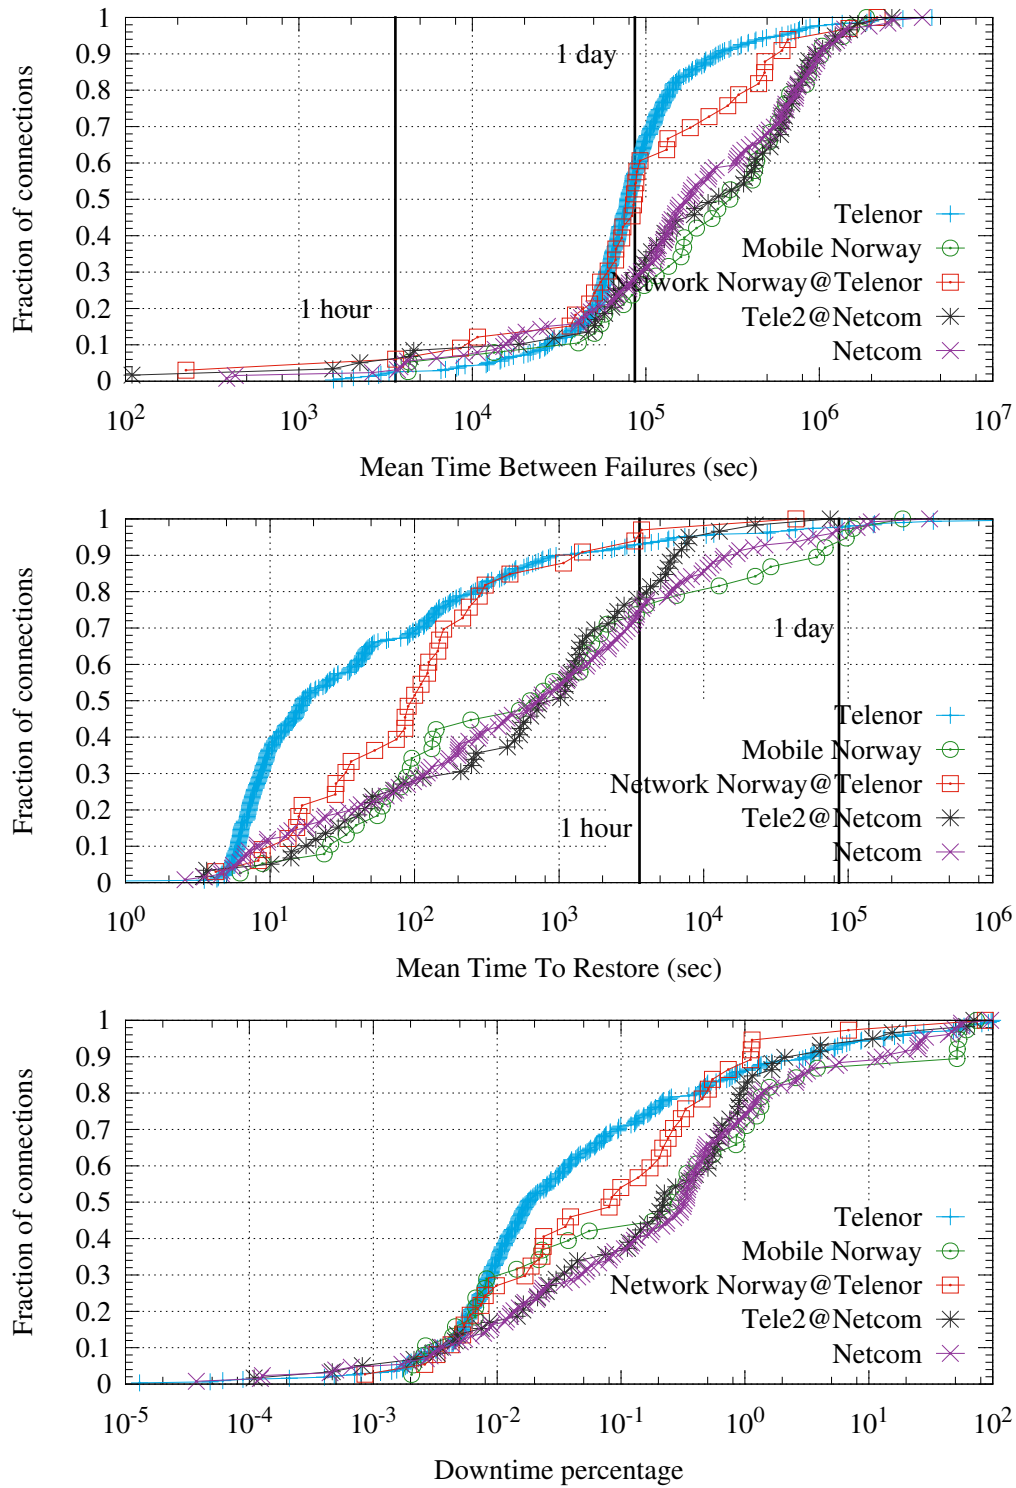


Figure 6.1: The statistics of connection failures.

6.1.2 Analysing connection failures

The stability of the tunnel that connects the UE to the CN depends largely on the RAN. Hence, to capture the effect of the radio access, we group our connections based on their respective RANs. Recall that the measured four UMTS operators use three RANs as illustrated in Figure 5.1. This gives us five logical networks in total, which are Telenor, Netcom, Mobile Norway (which includes Network Norway and Tele2 connections that use Mobile Norway's RAN), Network Norway@Telenor (which includes Network Norway connections that camp on Telenor's RAN), and finally Tele2@Netcom (which includes Tele2 connections that camp on Netcom's RAN). We use the camping information we collect from the modems to identify the connections that belong to the last three logical networks. For example, we classify a Network Norway connection as Network Norway@Telenor if it spends more than half of the time camping on Telenor's RAN, otherwise we classify it as Mobile Norway.

The three plots in Fig. 6.1 show the cumulative distribution function of the Mean Time Between Failures (MTBF), Mean Time To Restore (MTTR), and downtime percentage (due to PDP failures) for each connection in our data set, grouped by the five logical networks. We record distinct differences between operators, and observe a strong dependency between connection stability and the RAN. The statistics of Telenor connections and Network Norway@Telenor connections resemble each other. The same is true for Netcom connections and Tele2@Netcom connections. Despite that Mobile Norway is Network Norway's home RAN, the statistics of Network Norway@Telenor clearly differs from Mobile Norway's. The same is true for Tele2 and Mobile Norway, albeit to a lesser extent. This confirms the dominating role of the RAN in determining connection stability.

Differences between operators

Telenor and Network Norway@Telenor connections are less stable compared to the other three operators. About half of Telenor and Network Norway@Telenor connections fail at least once every day. For the other three operators this is the case for between one fourth (Mobile Norway) to one third of connections (Tele2@Netcom and Netcom). Telenor and Network Norway@Telenor, however, have much shorter MTTR compared to the other networks. Only 19% and 20% of Telenor and Network Norway@Telenor connections respectively have MTTR more than five minutes. The same numbers jump to 54% for Mobile Norway, 57% for Netcom, and 64% for Tele2@Netcom. These differences suggest that

the MTTR values for Netcom, Tele2@Netcom and Mobile Norway connections are influenced by a set of long lasting failures. To investigate whether these failures are the main factor behind the observed differences, we compute the median time to repair for all connections. While the median values are naturally smaller than the mean, the differences between operators remain consistent. For example, less than 9% of Telenor connections have a median time to repair longer than one minute compared to 33% for Netcom. Note that there are also slight differences, especially in the MTTR, between Network Norway@Telenor and Telenor. These differences can be attributed to the fact that many Network Norway@Telenor connections, though mainly camping on Telenor's RAN, spend some time camping on their home network as well. There are similar differences between Tele2@Netcom and Netcom but less pronounced. To check whether the observed difference between operators stem from varying coverage levels we measure the average RSSI for all connections. Figure 6.2 shows the Cumulative Distribution Function (CDF) of mean RSSI for each connection in all operators. All curves collapse onto each other indicating no systematic differences between operators. The same is true also for E_c/I_o .

Failure properties

Telenor and Network Norway@Telenor are dominated by frequent but short-lived failures compared to the other three networks. About half of Telenor and Network Norway@Telenor connections have MTTR less than 17 seconds and 90 seconds respectively. Looking closer at these short failures, we find that they are related to the RRC state of the connection, and they happen when the connection fails to be promoted from a shared channel (CELL_FACH) to a dedicated channel (CELL_DCH). This triggers the modem to reset the connection. As we demonstrate in Section 6.5, these short failures can have a drastic impact on applications performance. Netcom and Tele2@Netcom, on the other hand, have more long-lived failures that last for tens of minutes or even up to several hours. To gain a better insight into these long lasting failures, we investigate 157 failures in 27 distinct Tele2@Netcom connections which lasted for more than 1 hour. These connections are from NNE nodes that also have both a Netcom connection and a Telenor connection. Almost half of these failures (48.4%) affected the corresponding Netcom connections at the same time. The Telenor connections, however, remained stable. Hence, we believe that these long-lasting failures are not artefacts of our measurements. They seem rather related to the radio access availability, coverage,

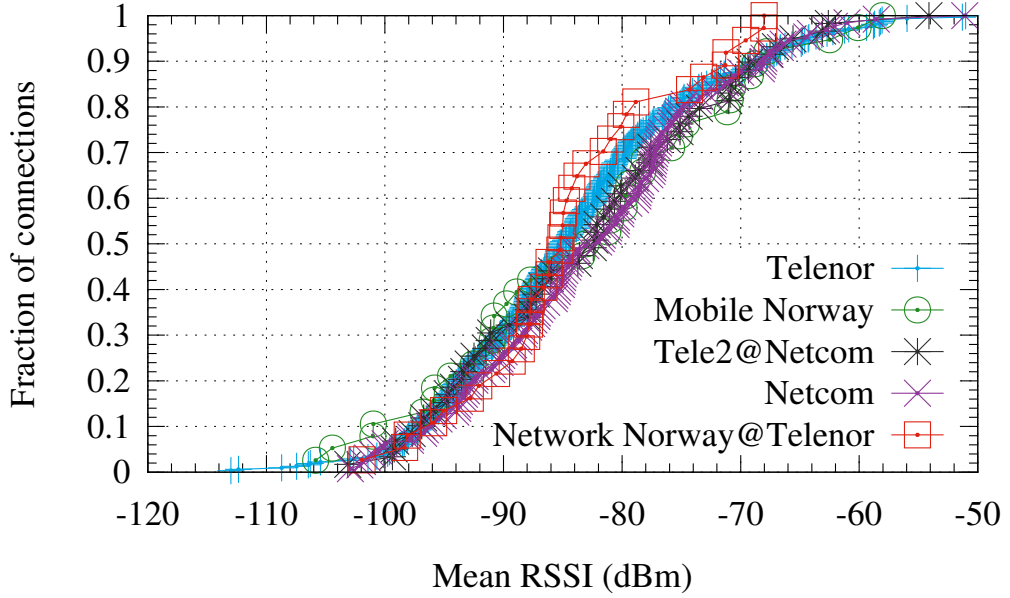


Figure 6.2: The CDF of the average RSSI per operator.

and possibly the interaction between the modems and the network.

Downtime

Telenor and Network Norway@Telenor connections have less overall downtime compared to the other networks. The percentage of connections experiencing more than 10 minutes of downtime per day ranges from 38% for Tele2@Netcom to 15% for Network Norway@Telenor. Failures that last more than 10 minutes are between 5.1% and 13.5% of all failures depending on the operator. They are, however, responsible for between 96.4% and 98.7% of the overall downtime. Besides characterising the overall connection downtime, we also investigate how connection stability has varied during our study period. To this end, we calculate the median daily downtime percentage per network measured as the median downtime across all connections available on each day. Figure 6.3 shows the time series of this metric for all networks throughout the study period. For all networks, the median daily downtime remains stable hinting at no significant changes in connection stability during our measurement period. Further, the time series are in line with the observations we made earlier in this section. Networks that share the same RAN exhibit similar median daily downtime. Also, Telenor and Net-

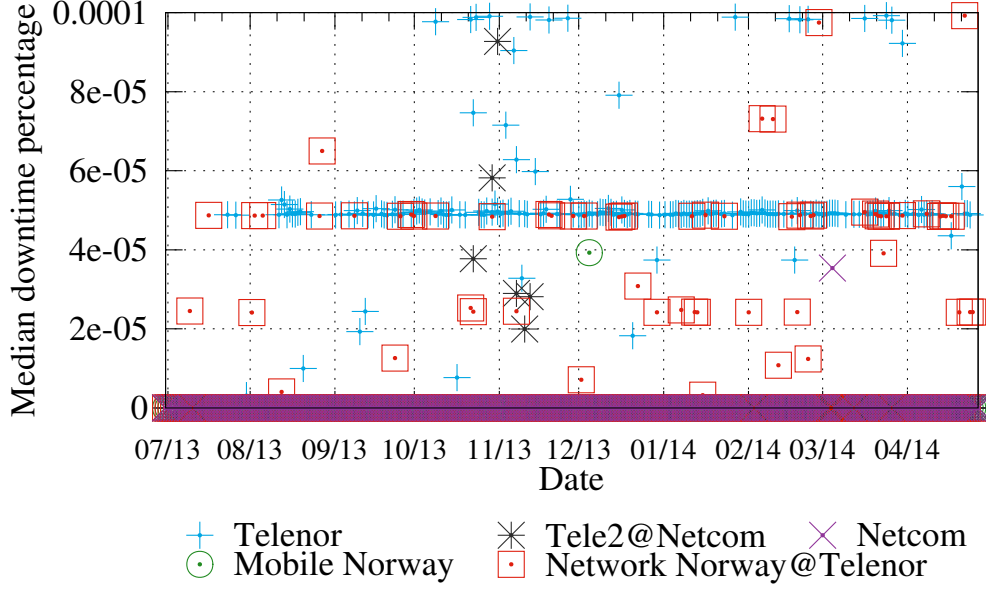


Figure 6.3: The daily median downtime percentage for each MBB operator.

work Norway@Telenor are characterised by a frequently observed median daily downtime of $5e-5\%$, which corresponds to a single outage of 4.32 seconds. This higher downtime percentage for both networks is consistent with our observation that they suffer more frequent short-lived failures compared to the other networks.

6.1.3 Correlating with metadata

To understand what may trigger the connection to be broken, we correlate the downtime due to PDP failures in a certain hour with the connection RAT (e.g., 2G or 3G), the average RSSI, and the average E_c/I_o in that hour. To correlate with the RAT, we say that a connection has 3G (2G) RAT in a given hour if it stays in 3G (2G) RAT for at least 70% of its available time. Further, to construct a meaningful correlation with the signal quality, we group the average RSSI values into five categories that correspond to the standard mobile phones signal bars: 1 bar (-103 dBm or lower), 2 bars (-98 dBm to -102 dBm), 3 bars (-87 dBm to -97 dBm), 4 bars (-78 dBm to -86 dBm), and 5 bars (-77 dBm or higher). We also group the average E_c/I_o values into three commonly used categories: Good ($0 \text{ dB} > E_c/I_o > -8 \text{ dB}$), Medium ($-8 \text{ dB} > E_c/I_o > -15 \text{ dB}$), Bad ($-15 \text{ dB} > E_c/I_o > -33 \text{ dB}$) [41]. We note that the RSSI measures the received signal strength, which includes all received compon-

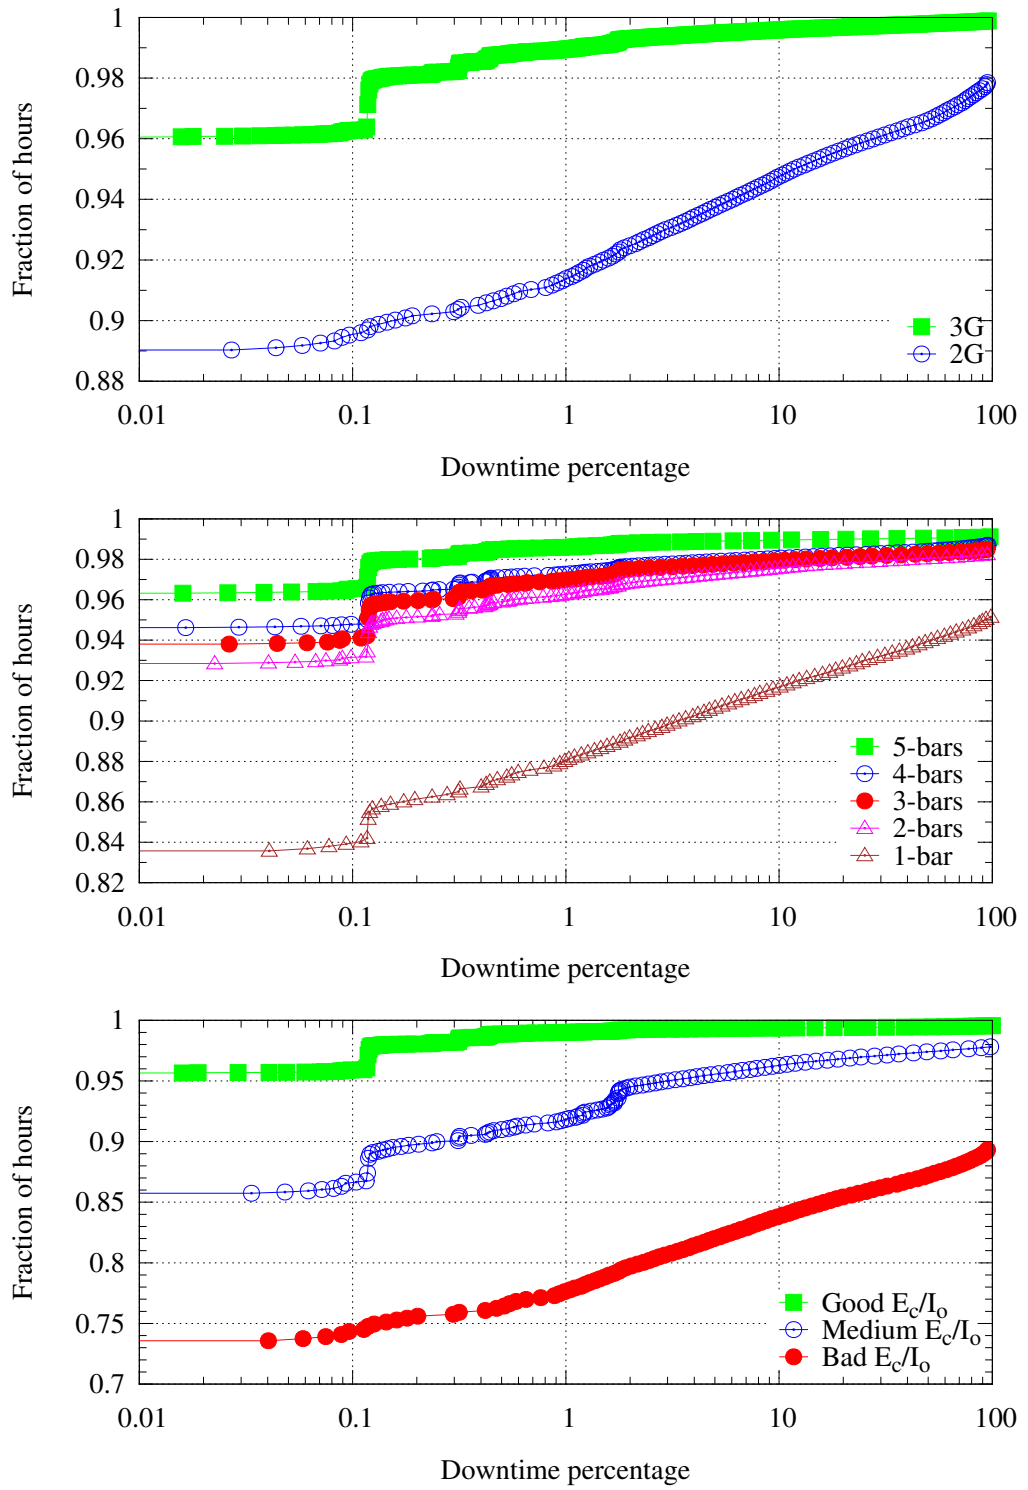


Figure 6.4: Downtime correlation with connection RAT, RSSI and E_c/I_0 .

ents (i.e. signal and noise). Hence, a high RSSI does not necessarily translate into a good radio condition. E_c/I_o on the other hand measures the signal quality (i.e. the signal to noise ratio) capturing both interference with ambient and surrounding noise as well as interference from cross traffic at adjacent frequencies.

The top panel in Figure 6.4 shows the CDF of the downtime percentage per hour, split according to the connection RAT. This plot includes all connections from all operators. The middle and bottom panels show the CDF of the downtime experienced at different RSSI levels and at different E_c/I_o levels respectively. The downtime is markedly higher for 2G connections than for 3G connections. We record 1% downtime or more in 1% of the hours with 3G connectivity compared to 9% of the hours with 2G connectivity. Further, as expected, downtime is influenced by the signal quality. Connections with an average RSSI equivalent to one signal bar have significantly higher downtimes. Beyond that, the differences between signal bar levels are less significant. Downtime also correlates strongly with E_c/I_o categories. We further observe that a sizeable fraction of hours with downtime exists even when having good E_c/I_o or RSSI. For example, we experience downtime in 5% of the hours characterised by a good E_c/I_o . This indicates that radio quality can not alone explain all connection failures. The stronger correlation between downtime and E_c/I_o as opposed to RSSI could be explained by the fact some locations of our nodes have relatively good signal but experience interference. We further correlate these three parameters to investigate whether a single measure is sufficient to describe the radio condition and consequently connection stability. Across operators, we do not observe clear correlation between RSSI and connection RAT. Poor E_c/I_o , however, strongly correlates with RSSI of one bar as well as with 2G connectivity. This suggests that E_c/I_o can be picked as a predictor of connection stability.

The above explained correlations are evident for all operators, but the relation between downtime and metadata is not always linear. For instance, the correlation between different E_c/I_o categories and connection stability is more evident in Telenor and Network Norway@Telenor than in Netcom and Tele2@Netcom. This suggests that disconnects in Telenor and Network Norway@Telenor are more often caused by the radio conditions, matching well with the short MTTRs discussed above. While such failures also exist for Netcom and Tele2@Netcom, they are masked by the dominating long-lasting failures.

6.1.4 Summary of findings

The results presented in this section show that many connections have a downtime that can be problematic for critical applications, such as alarms or payment systems. 15-38% of connections are unavailable more than 10 minutes per day on average. There are also clear differences in connection stability between operators. While Telenor experiences relatively frequent but short-lived failures caused by the failure to acquire a dedicated radio channel, other operators have less frequent but longer lived failures giving a higher overall downtime. We further find that the connection level reliability is highly dependent on the RAN. In particular there is a high correlation between downtime and the signal-to-noise ratio of the connection. Still, there is also a significant number of connection failures that can not be explained by radio conditions. These may be caused by congestion or central failures in the network.

6.2 Data plane: measuring packet loss

6.2.1 Measurement description and scenarios

This section looks at the networks' ability to deliver uninterrupted packet forwarding with an acceptable loss rate. Based on continuous end-to-end probing, we report on packet loss and on the duration of unavailable periods where no packets come through. We look at how these metrics are related to RAT, RRC state and signal quality. Finally, we identify particular events where a large fraction of connections simultaneously experienced abnormal packet loss rates.

Data plane reliability is measured by sending one 20 byte UDP packet to an echo server every second, and recording the reply packet from the server. A timestamp and an incremental sequence number is included in the packet payload for duplicate detection and round-trip time calculation. While the round-trip time normally is in the order of tens of milliseconds, we sometimes observe delays in the order of several seconds. Such high delays can sometimes be caused by excessive buffering [48]. We consider a packet to be lost if we do not receive a reply within 60 seconds.

This measurement test starts automatically on all network interfaces as they become active, and keeps running as long as the interface is up. When the connection is not available, the measurement script stops sending request packets and waits until the connection becomes available again. The measurement duration

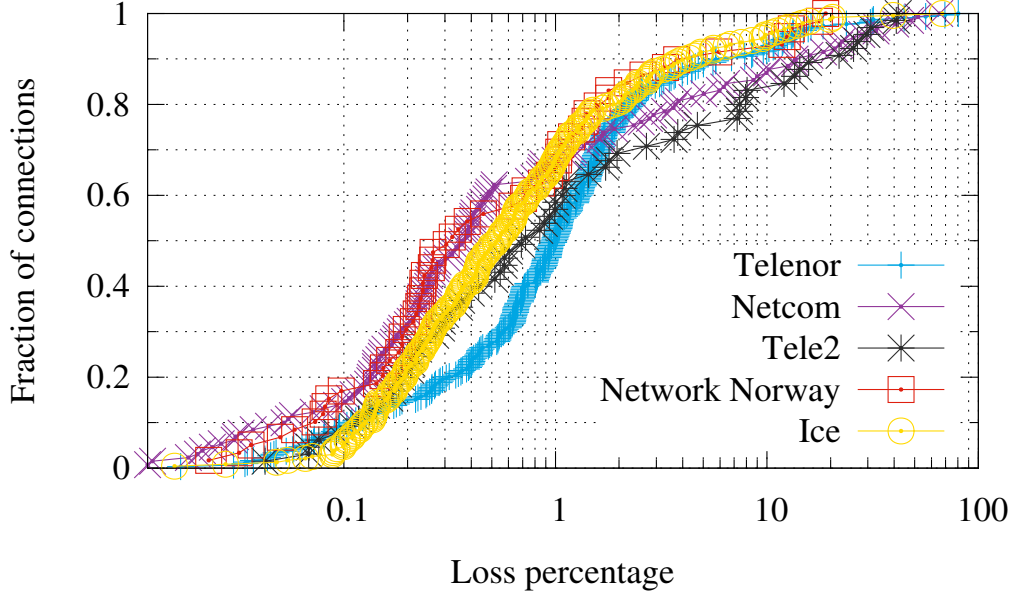


Figure 6.5: Loss rate for each MBB operator.

for each connection varies depending on how long the node was available and had working connections. In the following analysis, we require that we have at least 10 days of measurements to include a connection in the discussion.

6.2.2 Loss rate and loss runs

Overall loss rates

The CDF in Figure 6.5 shows the overall average loss rate for each connection in all operators. The loss rate is defined as (lost packets)/(sent packets) for the whole measurement period. Loss is relatively limited in all networks, and 50% of connections have less than 1% packet loss in all operators. We observe that relatively fewer Telenor connections have a very low packet loss rate compared to the other networks. 72% of Telenor connections have a loss rate higher than 0.5%, whereas this ratio is between 42 and 56% for the other networks. Telenor does not, however, have many connections with a high loss rate. Only 10% of connections have more than 5% loss, compared to 20% for Netcom and 23% for Tele2. Overall, Network Norway and Ice have the lowest packet loss rates.

There are diurnal patterns in packet loss in all networks, with higher loss rates

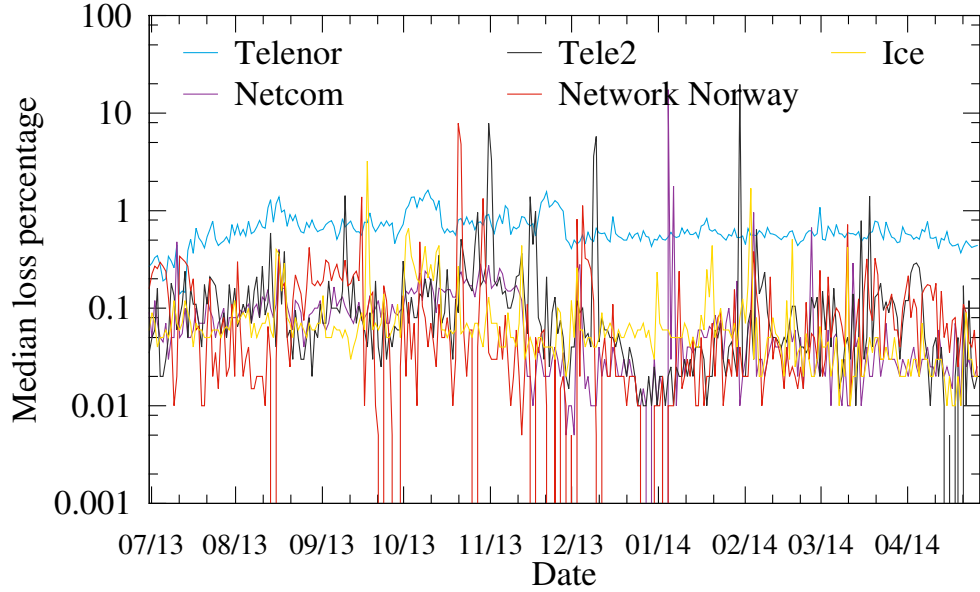


Figure 6.6: Median daily loss rate for each MBB operator.

in office hours when traffic is more intense. Ice, which has a larger fraction of home users, reaches their highest loss rates around 8PM. Packet loss in Telenor is consistently higher than in other networks throughout the day, also outside peak hours. We therefore believe that this higher packet loss is due to the RAN configuration rather than capacity limitations. To account for possible hardware, software or configuration changes over time in MBB networks, in Figure 6.6 we plot the median daily loss rate for each MBB operator. We see that during the whole measurement period the median loss percentage remains stable and conforms with the results show in Figure 6.5. Days with unusually high loss percentage are due to large failure events presented in Section 6.2.3.

Networks differ in the thresholds they use to promote a connection from CELL_FACH to CELL_DCH. Telenor is more conservative than the other networks in such promotions (see Section 6.3.1), and hence the measured Telenor connections are more often in CELL_FACH. For instance, a Telenor connection spends on average 65% of its time in CELL_FACH compared to 34% for a Netcom connection. Unlike in the other networks, Telenor connections have a higher loss rate in CELL_FACH than in CELL_DCH, indicating that this channel is operating closer

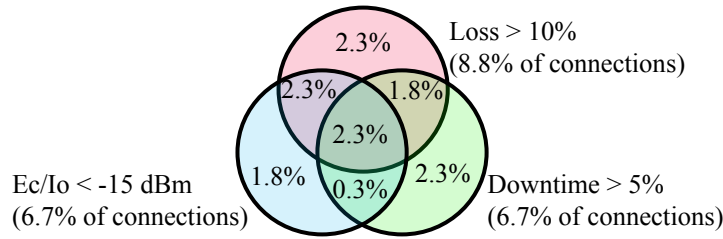


Figure 6.7: Loss, downtime and signal quality.

to its capacity limit². Telenor does, however, have higher packet loss than other operators also when on CELL_DCH, so this can not alone explain the difference.

To further explain the observed differences, we have looked at loss rates combined with the metadata collected in parallel with the measurements. We first observe that loss rates are similar for 2G and 3G connections in all networks, with the exception of Netcom, where 2G loss is higher. A typical Netcom connection experiences more than 1% packet loss in 60% (25%) of all hours when it is on 2G (3G). Loss increases in periods when connections perform vertical handovers between 2G and 3G/3G.

Not surprisingly, loss is higher in connections with poor signal quality. The Venn diagram in Figure 6.7 shows that many of the connections that experience a high loss rate also experience much downtime and have low E_c/I_o values. The diagram denotes the fraction of connections for which episodes with downtime, packet loss or low E_c/I_o values happened individually and alone (non-intersecting parts of the circles), as well as simultaneously (intersecting parts of the circles). Out of the 341 connections where we have all the necessary metadata to make this comparison, 8.8% have an average loss rate higher than 10%. As seen in the figure, most of these (73%) have either a downtime ratio >5%, $E_c/I_o < -15$ dBm, or both. We can not identify any systematic differences between operators in the relation between loss, downtime and signal quality.

Loss runs

From our measurement data, we identify each sequence of packets that were lost in a row. We call such sequence a *loss run*. Loss runs are important for how packet loss will affect the user experience. Since we send one packet every second,

²Telenor has previously confirmed to us that they have had capacity issues in the FACH channel.

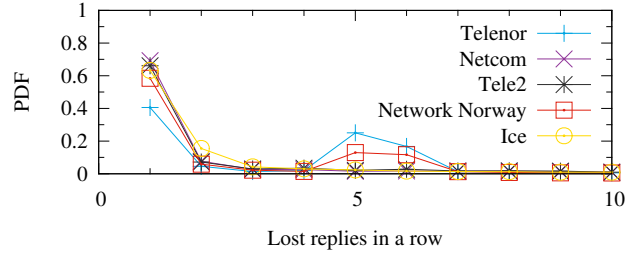


Figure 6.8: The distribution of loss run sizes across operators.

the size of a loss run is approximately equal to the number of seconds when no downlink traffic is received.

The distribution of loss run sizes is shown in Figure 6.8. Not surprisingly, the dominating size of a loss run is just one packet, accounting for 60% of loss runs for three of the networks. Telenor and Network Norway, however, also have many loss runs of size 5 and 6. As explained in Chapter 5, Network Norway connections sometimes camp on Telenor’s RAN, and we have confirmed that this is the case for connections with many loss runs of size 5 or 6. There is no clear periodicity in occurrence of loss runs of size 5 or 6, nor any dependence on signal quality, time of day or geographical area. Looking at the connection state during such loss runs, we find that they are normally caused by an RRC state demotion from CELL_FACH to IDLE. Re-establishing the CELL_FACH radio state and resuming packet forwarding takes 5-6 seconds, which is longer than an ordinary promotion [76]. While we do not know the exact cause of these demotions, we note that demotions can some times be caused by a network-initiated revocation of the radio access bearer. Such revocations can be triggered when capacity needs to be made available for other connections.

6.2.3 Capturing large failure events

Next, we discuss *large failure events*, where many connections in an operator experiences abnormal packet loss at the same time. Such events will normally be caused by failures in the CN. To identify and analyse large failure events, we first divide our measurement time series into 5 minute intervals, and calculate the loss rate for each connection in each interval. Additionally, to account for downtime periods, all 5 minute intervals when a connection was unavailable are assigned a 100% loss rate. We consider that a connection has an abnormal loss rate in a par-

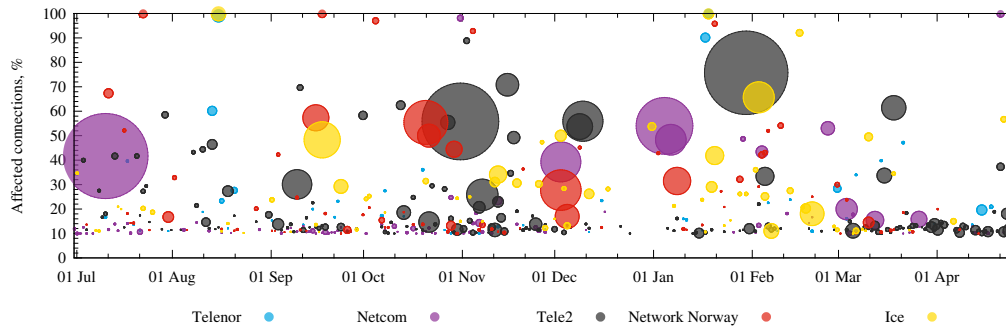


Figure 6.9: Large failure events July 2013 – May 2014.

ticular 5 minute interval if more than 10% of the packets are lost. A large failure event is defined as a period of one or more intervals when more than 10% of all connections in an operator has abnormal packet loss.

Figure 6.9 shows a visual representation of all large failure events recorded during our measurement period. Each event is represented by a circle (multiple events in the same day are merged). The diameter of the circle reflects the severity of the event, and can be thought of as the total volume of lost traffic. This is calculated as the product of the fraction of affected connections, the average loss rate in the affected connections, and the duration of the lossy period. The fraction of affected connections is also represented on the y-axis.

As seen in the figure, networks experience large failure events with varying severity and frequency. Short-lived events with limited loss and affecting 10-20% of connections happen on a weekly basis in all networks. These events might be attributed to short-lived congestion, and may be considered part of normal operation. There are, however, also a number of larger failure events that can severely influence the user experience.

The collected measurements can normally give a good idea about the underlying cause for the more severe large failure events. By looking at the geographical distribution of the affected connections, the affected RAN(s), loss intensity and other parameters, we can often pin the root cause to either the transmission network, the CN, or to the interconnection between operators.

For example, on Sept. 10 2013, most Tele2 connections experienced 20-50% packet loss for around 40 minutes. Similar behaviour repeated itself on Oct. 28 and on Nov. 16. These failures happened outside of maintenance windows, and affected connections from all parts of the country, and were therefore likely re-

lated to a component in the CN. Tele2 has later informed us that these events were probably caused by a failure in a load balancer in their CN.

One of the largest failure events recorded in our measurement period took place in Tele2 on Nov. 1-2, when 86% of connections were unavailable for more than 6 hours. 41% of the affected connections lost the PDP context during this period, while the others maintained a valid IP address but could not send or receive any data. The failure affected only Tele2 connections camping on Netcom's RAN, and Tele2 has confirmed that the root cause of the failure was a failing component in the peering between these two networks.

An interesting event took place on Oct. 21-22, and affected all Network Norway connections camping on Telenor's RAN. These connections experienced around 50% packet loss for more than 11 hours. We also observe that the packet loss rate was higher at times when traffic volumes are higher, indicating a capacity problem in the connection between these two networks. Network Norway later confirmed that this event was caused by a routing problem that sent traffic through a roaming exchange point normally used for international roaming traffic only. This link is dimensioned for significantly lower traffic volumes, and could therefore not support the offered load.

In our discussions with the network operators, we have learned that several of the loss events identified in this study were not detected by their internal monitoring systems. These systems detect situations where many customers lose their Internet connections, but not necessarily situations where packet loss is less than 100%. This shows the potential of end-user measurements, and illustrates how they can help operators discover weaknesses and failures in their networks.

6.3 Data plane: dissecting packet loss

In this section, we investigate whether we can quantify and explain loss using end-to-end measurements. This study is based on the measurement data from June 2013 to July 2014, when the majority of our connections were on 3G. The study is therefore limited to 3G connections only.

We leverage connections' metadata to group connections from the same operator in order to gain insights into the root causes of packet loss. Specifically, we group connections by their serving cell and serving RNC, and identify loss that affects several connections in one group at the same time. Our results give several invaluable insights into the nature and causes of loss in MBB networks.

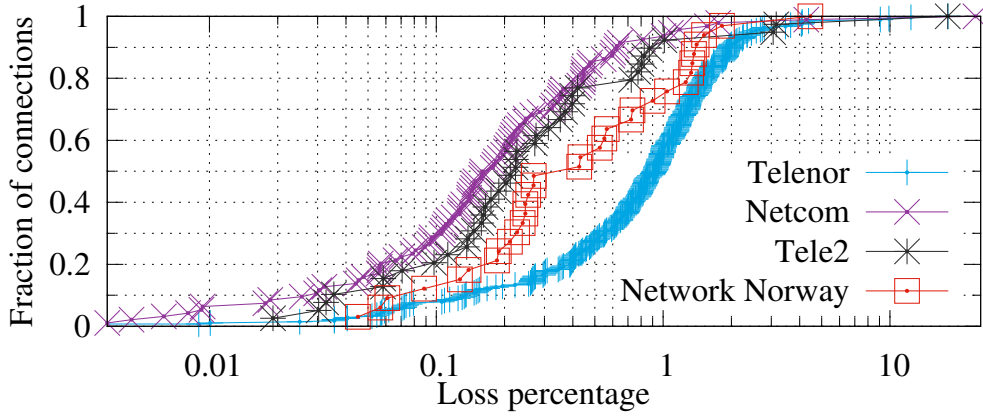


Figure 6.10: Loss rate for each MBB operator. We see clear differences between operators.

6.3.1 Packet loss statistics

This section looks at the overall statistics and characteristics of measured packet loss. Using continuous end-to-end probing and rich connection metadata, we report on per-connection packet loss. We further investigate the interplay between radio channel types and radio resource control on the one hand and packet loss on the other.

In the following analysis, we only include connections that contribute at least one week of measurements. To avoid any artefacts that could be caused by poor coverage, we only include connections that have an average signal strength greater than -102 dBm. Further, as some of our connections experience mode switches from 3G to 2G and vice versa, leading to excessive packet loss, we only include data from periods when connections were predominantly on 3G. Additionally, we filter out maintenance window periods and cases in which we had problems with our infrastructure. Lastly, in our analysis we consider only the data when the measurement script was able to send 300 packets via a single MBB connection during a five-minute interval. We do this to avoid introducing packet loss, which might happen right after the PDP context is established or before abnormal termination.

Loss rate

Figure 6.10 shows the CDF of overall loss rate per connection for all operators. This loss rates share similar patterns with ones shown in Figure 6.5, but are gener-

ally lower due to more extensive filtering of five-minute bins as described above. As seen from the figure, loss is relatively small in all networks; at least 50% of connections experience less than 1% packet loss across operators. We also see clear differences between operators. 72% of Telenor connections suffer more than 0.5% loss, whereas this ratio is between 15% and 43% for the other networks. As for the fraction of connections with excessive packet loss (e.g. a few percent), we do not measure clear differences between networks.

Loss patterns in Netcom and Tele2 show clear similarities, and half of Tele2 and Network Norway connections show comparable loss. Furthermore, the worst 20% of Telenor and Network Norway connections exhibit similar loss. These observations hint at different loss components that lie both in the RAN and beyond it. Recall that Tele2 and Network Norway own the same RAN, *Mobile Norway*, which covers only some parts of the country. Outside areas covered by their RAN, Tele2 and Network Norway use Netcom's and Telenor's RAN respectively.

The role of RRC

Radio states differ in available resources and consequently performance, hence it is important to control for their effect when analysing packet loss. Our connections are either on CELL_FACH or CELL_DCH, since they continuously send data. Connections that only send 20-byte UDP pings remain on CELL_FACH. Note that the connections we use for remotely managing measurement nodes and for transferring measurement logs are typically on CELL_DCH. To understand the effects of RRC state on packet loss, we divide our measurement period into consecutive five-minute bins. We then classify each connection bin into one of four categories based on the recorded RRC states within it:

1. **DCH.** A connection bin is classified as DCH if the connection was on CELL_DCH throughout the bin.
2. **FACH.** A connection bin is classified as FACH if the connection was on CELL_FACH throughout the bin.
3. **Mixed-UP.** These are connection bins in which connections experience either state promotion from CELL_FACH to CELL_DCH or a demotion in the opposite direction.
4. **Mixed-Downgraded.** These are connection bins in which connections experience a state demotion to CELL_PCH or IDLE. Note that this should not

happen since our connections are always sending data.

Table 6.1: Time spent (TS) and loss induced (LI) by each network on different RRC categories.

Net.	TS LI	FACH	DCH	Mixed-UP	Mixed-Downgraded
Telenor	TS	56.3%	4.4%	21%	18.3%
	LI	4.7%	1.9%	9.6%	83.7%
Netcom	TS	26.8%	59.2%	13.3%	0.7%
	LI	9.4%	35.9%	19.2%	35.5%
Tele2	TS	35.2%	60.4%	2.3%	2.2%
	LI	12.1%	32%	2.3%	53.6%
Network Norway	TS	64.2%	5.2%	19.2%	11.5%
	LI	6.8%	0.8%	6.3%	86%

After classifying all connection bins, we calculate the fraction of time spent and loss induced on different RRC categories for each operator. Specifically, we sum all connection bins and lost packets in them from the same operator, which gives the total number of measurement bins we collected and the total number of lost packets for this operator. We then calculate the number of these bins that are in different RRC categories. Table 6.1 presents the results of these calculations. The most important observation is that all operators spend some time in the Mixed-Downgraded state, which reflects unexpected pathological behaviour. It also shows that most loss happens while on Mixed-Downgraded for all networks except Netcom, where loss induced in DCH and Mixed-Downgraded is similar. According to the 3GPP standards, as long connections are actively engaged in data exchange, they should never be demoted to IDLE or CELL_PCH. We note that these transitions happen to connections on DCH as well as to connections on FACH. The presence of such demotions indicates that they are not strictly inactivity based. This pathological behaviour is most evident in Telenor and Network Norway. Recall that Network Norway connections use Telenor’s RAN in locations not covered by their RAN. This indicates that the observed pathological behaviour is RAN-related, which is consistent with the fact that RRC states are managed by RNCs.

Overall, connections that use Telenor’s RAN undergo markedly more state transitions, including both pathological and standard promotions and demotions

compared to connections that use other RANs. Additionally, the percentage of time spent on FACH and DCH also varies widely across operators. This is mainly an artefact of the way NNE connections are used for tasks other than sending the 20-byte pings. For instance, in nodes with UMTS connectivity only, logs and measurement data are transferred over Netcom and Tele2 connections, which explains why these two operators spend much more time on DCH.

Knowing the RRC categories in which our connections spend the most time and which carry the most loss, we now move to measuring the loss rate of individual connections for each category. For each connection, the loss rate while in category C_i is defined as (the number of lost packets while on C_i)/(the number of sent packets while on C_i). The four panels in Figure 6.11 show the loss rate in each category for all four operators. Across the board, Mixed-Downgraded stands out as the lossiest category, in accordance with the induced loss numbers in Table 6.1. This is expected since Mixed-downgraded corresponds to cases in which connections are demoted to IDLE or CELL_PCH. In these states connections lose their data connection to the network and hence are not able to exchange data. Loss in periods with regular state transitions (i.e. Mixed-UP) is the second worst. In such five-minute bins, a connection spends a fraction of the time in FACH, another fraction in DCH, and a short period transiting between them. We expect loss during the combination of FACH and DCH states to be comparable to the loss in FACH and DCH categories separately. Hence, we believe that the difference in loss between Mixed-UP periods and FACH (DCH) periods is mainly due to transitions between states. Investigating this further with one of the four operators studied confirmed that. Loss can happen during transitions because the UTRAN RLC, a protocol responsible for interfacing between upper layers and the MAC layer [41], has different Service Data Unit (SDU) sizes for FACH and DCH. Interestingly, differences between FACH and DCH are mostly minor. Thus RLC SDUs in transit must be buffered when undergoing state transitions [77] and re-transmitted once they are complete. In fact, the mentioned operator enabled the retransmission as a result of our observations. Before that, the disabled retransmissions during the RLC re-establishment affected many applications using unreliable transport protocols, such as UDP, including our UDP ping measurement.

The breakdown of loss along RRC states gives more insights into the measured differences between operators (See Figure 6.10). Telenor exhibits more loss because it spent $\approx 40\%$ of the time in Mixed-Downgraded and Mixed-UP. Network Norway exhibits less loss, although it spent $\approx 30\%$ of the time in these two states. Looking closer at Network Norway, we find that the excessive loss in Mixed-

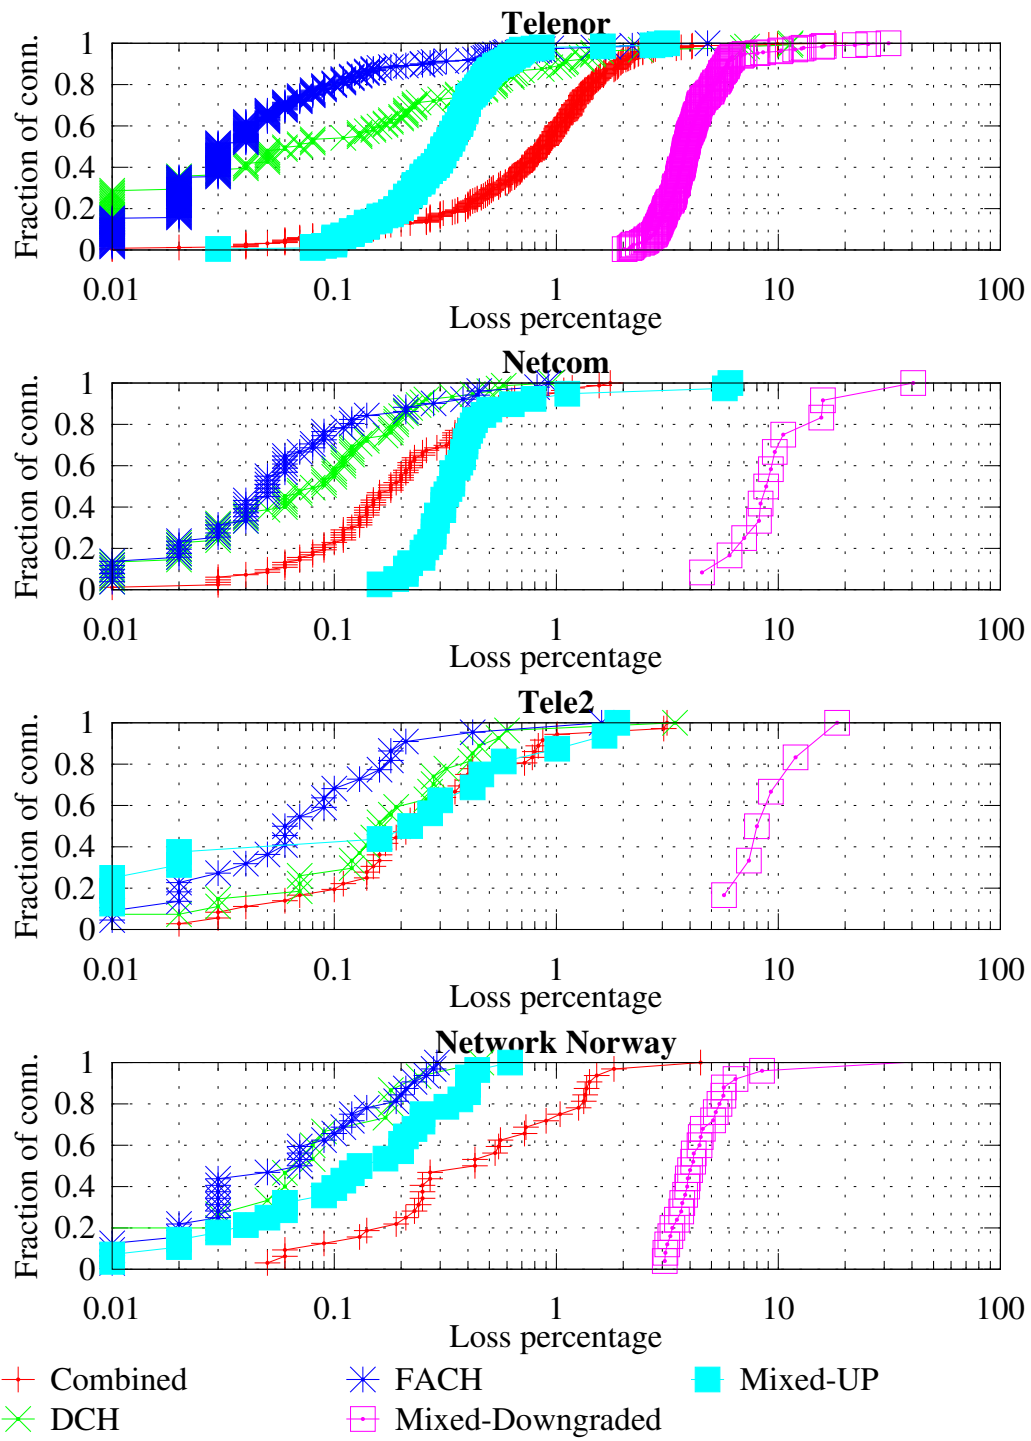


Figure 6.11: Loss rate experienced during different RRC bins. We record very high loss rate during periods with state transitions.

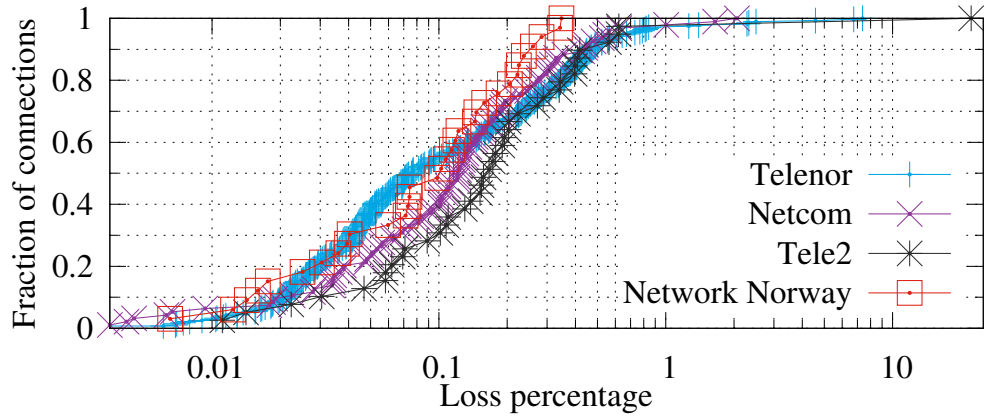


Figure 6.12: Loss rate for each MBB operator considering only non-pathological states.

Downgraded is offset by negligible loss in the other three categories. Figure 6.12 illustrates the CDF of the overall loss when considering only loss that happens in FACH, DCH, and Mixed-UP bins. Overall, loss decreases significantly, at most 9% of connections experience loss more than 0.5%. Further, differences between operators become less pronounced. Hence, avoiding these pathological transitions can greatly reduce loss.

These results inspired Telenor to revisit its configuration:

- enable inactivity based timer for the FACH to PCH transition (previously the transition was happening when maintaining a low bit rate);
- enable packet retransmission during RRC re-establishment;
- and few other changes on their RNCs.

This has resulted in a major reduction in pathological transitions and overall loss rate (see Section 6.3.7).

Summary of findings

Our results demonstrate clear differences in packet loss between operators, and show that most loss is caused by RRC state transitions, both pathological and standard. The former transitions force active connections into idle and paging modes, deallocating the radio resources they need for sending data. We expect differences between operators to diminish if pathological transitions are avoided.

6.3.2 Correlations in loss

In order to reduce packet loss, we first need to pinpoint its potential causes. To this end, we correlate loss measured for pairs of connections from the same operator. We perform these correlations at the operator level, RNC-level, and cell-level to determine whether observed loss is independent or caused by events that impact multiple connections simultaneously (e.g. network congestion and outages).

To correlate a pair of connections, we divide each connection UDP-ping time series into five-minute bins, and calculate the loss ratio in each bin. We use these time series to determine the likelihood that this pair of connections experiences loss equal to or greater than $X\%$ in the same five-minute bin. Let $P(A)$ denote the (empirical) probability that a connection A has a loss rate equal to or greater than $X\%$ in a given five-minute bin, and $P(B)$ denote the same for a connection B . We calculate the conditional probability $P(A|B)$ of A having loss given that B has loss and compare it to the unconditional probability $P(A)$. If the conditional probability ratio $R = P(A|B)/P(A)$ is close to 1, it means that connections A and B experience loss largely independently, while a high R means that they tend to experience loss together. Note that by Bayes' law, $P(A|B)/P(A) = P(B|A)/P(B)$ [51]. We have experimented with setting X to 1%, 3%, and 5% packet loss which mostly gave qualitatively similar results. In the rest of this section, we stick to the 3% threshold (i.e. the loss of at least nine packets in a 5-minute bin) to avoid any spurious correlations that $X = 1\%$ may cause as well as to avoid focusing only on rare loss events which limit us to 5% or more loss. The loss in bins with 3% or more packet loss is between 75% and 91% of the overall loss depending on the operator.

6.3.3 Operator-wide correlations

Thus far, our findings show that pathological state transitions are responsible for a significant fraction of the measured packet loss. It is unclear, however, whether this loss is correlated (e.g. these transitions happen because of a sub-optimal RNC configuration that affects several connections simultaneously). To verify this, we calculate the conditional probability ratio R for all pairs of connections from the same operator considering all measured five-minute bins. We also calculate the same ratio correlating only bins of certain RRC category. In other words, for each connection, we construct two additional time series. The first time series include all five-minute bins in which the connection was on Mixed-Downgraded,

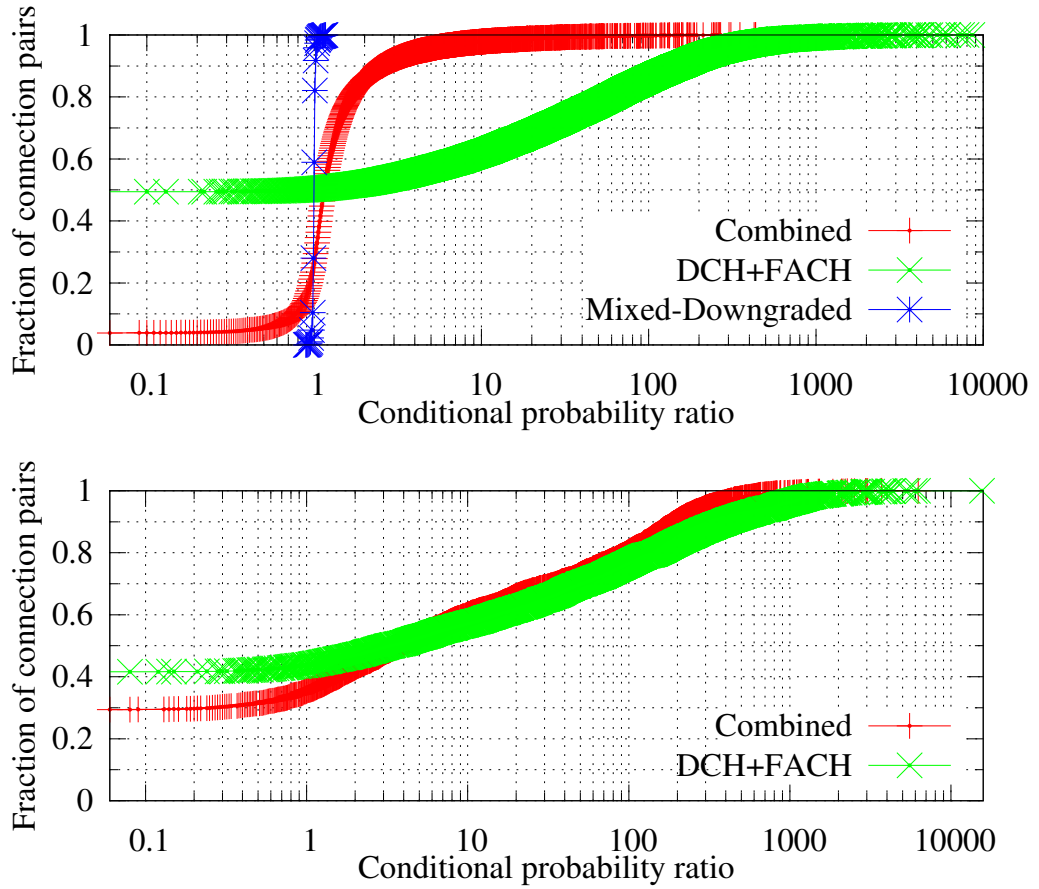


Figure 6.13: Correlations for loss $\geq 3\%$ in different RRC categories for Telenor's RAN (top) and Netcom's RAN (bottom). Loss in DCH and FACH exhibits relatively strong correlations.

the second involves bins of the three non-pathological types DCH, FACH, and Mixed-UP, which we refer to as DCH+FACH.

The plots in Figure 6.13 present the CDF of R for different RRC states for Telenor's RAN and Netcom's RAN. In the remainder of this section, we leave Mobile Norway RAN aside because of its small sample size. When correlating bins regardless of their RRC state (denoted as "Combined" in the figure), Telenor's RAN demonstrates little correlation with 90% of the connection pairs having $R \leq 2$. Netcom's RAN, however, appears more correlated, with only 40% of pairs having $R \leq 2$. Both RANs are, however, characterised by a lack of correlation for loss related to pathological state transitions. For Telenor's RAN, loss in Mixed-

downgraded is independent with $R \approx 1$. For Netcom’s RAN, there is not a single correlated pathological bin, which explains why Mixed-Downgraded is missing in the lower panel. Furthermore, DCH+FACH demonstrates strong correlations for both operators, with about 40% of connection pairs having $R \geq 10$. We note that for connection pairs in the FACH+DCH category with $R \leq 2$, the median common measurement period (i.e. the time in which both connections were present) is 60 days and 93 days for Telenor’s and Netcom’s RAN, respectively. These periods are long enough to observe rare network events and to avoid spurious correlations. Furthermore, R equals zero for about 50% (40%) of pairs in Telenor’s (Netcom’s) RAN meaning that there is not a single common lossy bin. Investigating these pairs shows that they have less than a third common measurements period compared to pairs with $R \leq 2$. Hence, this lack of correlation can be attributed to the interplay between the relatively short common measurement periods and the measured rarity of excessive loss in MBB in general. In other words, we need longer common periods, as we have for pairs with $R \leq 2$ to be able to identify correlated loss. Based on these observations, we note the following points. Loss during pathological RRC transitions impacts connections independently. Further, correlating loss bins that are characterised by different RRC states may lead to wrong conclusions. For example, loss in Telenor’s RAN appears uncorrelated if we do not split it according to the RRC state.

The strong correlation in loss during DCH+FACH suggests that the causes of this loss do not affect connections individually. Possibly, these causes are related to components in the MBB network that serve multiple connections simultaneously (e.g. NodeB, RNC, or GGSN). The remainder of this section investigates the role of such components by evaluating correlations in DCH+FACH loss at the RNC and cell levels.

6.3.4 RNC-level correlations

To be able to perform RNC-level correlations, we first need to map connections to RNCs. NNE collects a rich set of metadata about connections, including the LAC and the UTRAN CID. The LAC indicates a group of towers that are served by the same RNC. However, there is no one-to-one mapping between LACs and RNCs. The UTRAN CID is a 6-byte hexadecimal value, the first 2-bytes represent the serving RNC-ID, while the other 4-bytes represent the CID. We use this to determine the RNCs that serve Netcom’s RAN connections. Unfortunately, Telenor’s RAN sets the RNC-ID bits to the first digits of LAC, but not the real RNC

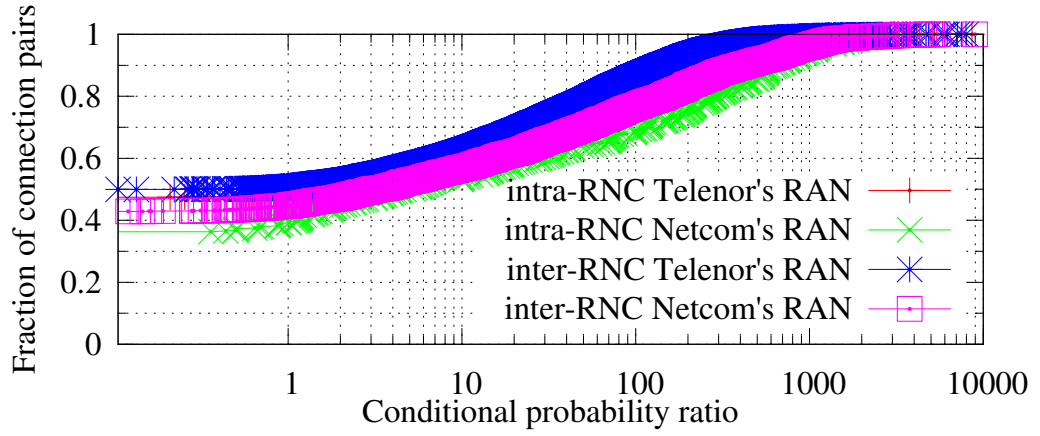


Figure 6.14: Correlation in loss for connections served by the same RNC vs connections served by different RNCs.

identifier. Hence, we resort to using an RNC coverage map provided by Telenor that indicates all RNCs and their geographic scope. Telenor's RAN comprises 12 RNCs, while Netcom's RAN comprises 14. The number of connections served by each RNC depends on the NNE nodes' deployment. The percentage of connections per RNC varies from 2.4% to 25.5% for Telenor's RAN and from 0.8% to 19.7% for Netcom's RAN. We have more connections associated with RNCs that serve the three largest cities in Norway.

We calculate the conditional probability ratio for all connection pairs that are served by the same RNC, which we refer to as the intra-RNC correlation. We only consider pairs that shared the same RNC for at least one week. In addition, we correlate pairs of connections that are not served simultaneously by the same RNC, which we refer to as inter-RNC correlation. Note that we are only correlating loss that happens in DCH+FACH. Figure 6.14 shows the CDF of inter-RNC and intra-RNC correlations for Telenor's and Netcom's RANs. For both RANs, differences between intra-RNC correlations and inter-RNC correlations are low. Note that correlations for both groups remain high, suggesting that the causes of these correlations lie beyond the RAN. Next, we investigate how many correlations at the RNC level can be attributed to correlations at the cell level.

6.3.5 Cell-level correlations

In order to study cell-level correlations, we need to have an adequate number of connection pairs that are served by the same cells simultaneously. Unfortunately, the NNE measurement infrastructure was not deployed such that it maximises the number of measurement nodes covered by the same cell towers. This makes the task of finding connections that share serving cells difficult. We scan our data set to identify connections that attach to the same cells simultaneously by looking at two types of nodes. First, nodes that are in the same cell tower ranges, thanks to NNE dense deployment in the three largest cities in Norway. Second, nodes that have two connections served by the same RAN (e.g. a Telenor connection and a Network Norway connection that camps on Telenor's RAN). These two connections are essentially served by the same cells and the same RNCs. We identify 34 and 31 cells belonging to Telenor's RAN and Netcom's RAN respectively that serve multiple connections simultaneously. Note that the overall number of distinct cells in our data set is 1988 for Telenor's RAN and 458 for Netcom's RAN.

We calculate the conditional probability ratios for all connection pairs that are served by the same cells, which we refer to as the intra-cell correlation. We only consider pairs that shared the same cell for at least one week. In addition, we correlate pairs of connections that are not served simultaneously by the same cells, which we refer to as the inter-cell correlation. Note that we are only correlating loss that happens in DCH+FACH. Figure 6.15 shows the CDF of inter-cell and intra-cell correlations for Telenor's RAN and Netcom's RAN. For both RANs, intra-cell correlations are mostly higher than inter-cell correlations. Inter-cell correlation is reasonably high ($R \geq 10$) for a sizeable fraction of pairs, but it is still smaller compared to the above presented operator-wise and inter-RAN correlations. We get $R = 0$ for about 80% of inter-cell pairs in Telenor's RAN and 65% of inter-cell pairs in Netcom's RAN. These numbers are higher than their network-wide and RNC counterparts. These discrepancies stem from the fact that our inter-cell and intra-cell pairs have very short common periods compared to network-wide pairs or RNC pairs. For instance, the median common measurement period for intra-cell pairs is 10 days for Telenor's RAN and 15 days for Netcom's RAN. A typical connection in our data set alternates between 3 or 4 cells, so a pair of connections in the range of the same group of cells may end up connecting to different cells within the group reducing the common measurement period. These periods are too short when measuring correlations in loss given the demonstrated rarity of excessive loss in MBB. To confirm this, we re-calculate the inter-cell and

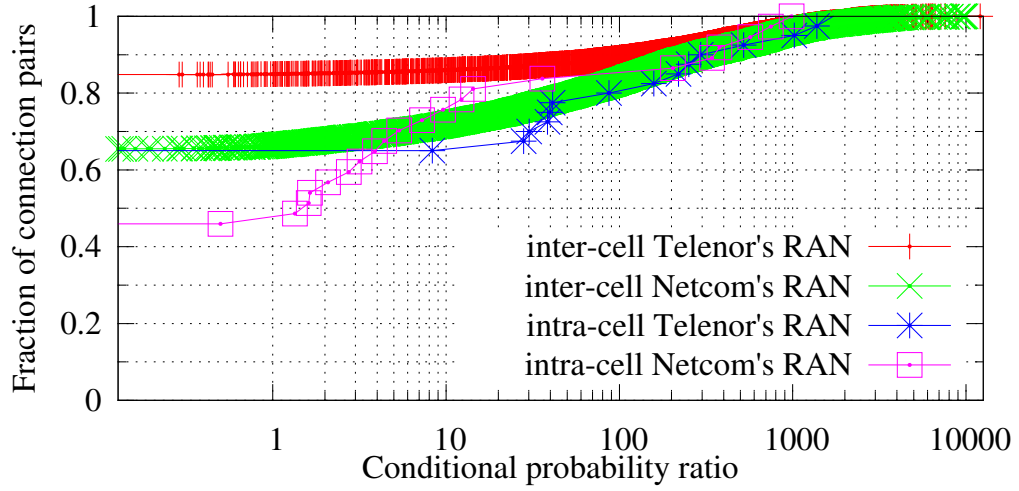


Figure 6.15: Correlation in loss for connections served by the same cell vs connections served by different cells.

intra-cell correlations considering five-minute bins with loss $\geq 1\%$ (i.e. greater likelihood of loss). We find that the differences between inter-cell and intra-cell correlations diminish, giving results similar to the RNC-level correlations.

Summary of findings. The results of this section show that loss experienced during the pathological Mixed-downgraded transition is mostly uncorrelated. Hence, these transitions reflect a behaviour that affects individual connections rather than several connections simultaneously. The loss experienced during DCH+FACH, however, exhibits a strong correlation. We measure a negligible difference in the likelihood of simultaneous loss between connection pairs served by the same RNC and those served by different RNCs. This suggests that the measured RNCs are not congested. The same is true at the cell-level if we control for the length of the common measurement period. Finally, we note that the presence of strong correlations between arbitrary connection pairs hints at causes that lie beyond the RAN.

6.3.6 Possible causes

Our results indicate that there is a non-trivial fraction of correlated loss. Further, the causes of these correlations seem to lie beyond the RAN. This section makes a one step closer to understanding potential causes of loss.

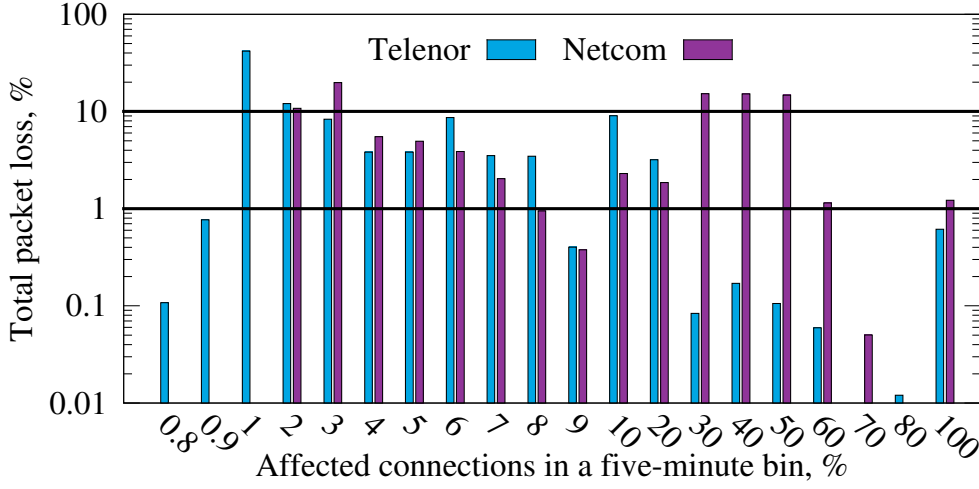


Figure 6.16: Loss rate quantification for Telenor and Netcom MBB networks.

Impact of the loss episodes

The pairwise correlations presented in the previous section only show whether a connection pair is likely to experience loss at the same time. It does not show how many connections in total experience loss simultaneously. The number and location of connections that suffer loss at the same time give an idea about where the root causes lie. For instance, if the number of impacted connections is high and these connections are served by different RNCs, then the root cause lies beyond the RAN. Hence, in order to identify the root causes of most loss, we need to quantify how much of the overall loss is related to loss episodes that impact a few connections as opposed to those impact a significant fraction of connections.

To measure that, for each connection we identify all five-minute bins with loss equal to or greater than 3%. Next, we merge bins from all connections, giving us a list of all five-minute bins where at least a single connection experiences 3% or more loss. Every bin is, therefore, characterised by two numbers: the number of affected connections, and the number of lost packets on all affected connections. We use these two numbers to quantify how many loss episodes with a given number of affected connections contribute to all packets in bins with loss $\geq 3\%$. We limit ourselves to bins with loss $\geq 3\%$ to avoid spurious correlations. Further, the majority of loss happens in these bins (i.e. between 71% and 91% of the overall loss depending on the RAN).

Figure 6.16 shows histograms of the percentage of packets lost, binned by the

percentage of connections impacted, for Telenor and Netcom. The two horizontal lines mark the 1% and 10% loss levels. We record clear differences between the two operators. Netcom's histogram exhibits several pronounced modes for the impact range between 30% and 60% of connections. One third of the loss happens in bins where 1/3 or more connections are affected. Telenor histogram also exhibits clear modes but they correspond to much smaller impact ranges. About half of the loss happens in bins where 2% or fewer connections are affected and $\approx 84\%$ of all loss happens in bins where 10% or fewer connections are affected. Furthermore, for both networks loss episodes that impact all connections simultaneously account for 1% of overall packet loss.

Our findings indicate that Netcom is more prone to loss episodes that affect a large fraction of connections simultaneously. This fraction (i.e. over one third) is evidently higher than the percentage of connections served by the most populous RNC in Netcom's RAN (19.7%). Hence, the root causes of this highly correlated loss lie in the CN or in the interconnection between Netcom and the Internet. In contrast, Telenor is less prone to episodes that impact a large fraction of connections simultaneously, suggesting that Telenor's core is better provisioned than Netcom's. Note that, despite the rarity of widely correlated loss in Telenor, a sizeable fraction of loss (16%) happens in bins where 10% or more geographically-diverse connections are affected.

Diurnal patterns in loss

An increase in traffic at certain times of the day congesting different network components is a possible cause for loss. To investigate this, we calculate average hourly loss per RRC category for each operator. To account for cases when connections became stale and hence 100% lossy for longer periods, we remove all 5-minute bins with 100% packet loss. Further, we filter out 5-minute bins where we see 1% or more packet loss affecting 50% or more of our connections from two or more MBB networks. We do this to rule out bias coming from our infrastructure or Internet service provider.

Figure 6.17 shows variations in the average hourly loss per RRC category normalised by the smallest value in the respective category. Regardless of the RRC category and the network, loss rates are higher during daylight hours, from around 6 in the morning until 22 in the evening. This hints that a certain fraction of loss is attributed by the increased traffic and therefore congestion during these hours.

We observe that in Telenor, both *Mixed-UP* and *Mixed-Downgraded* RRC cat-

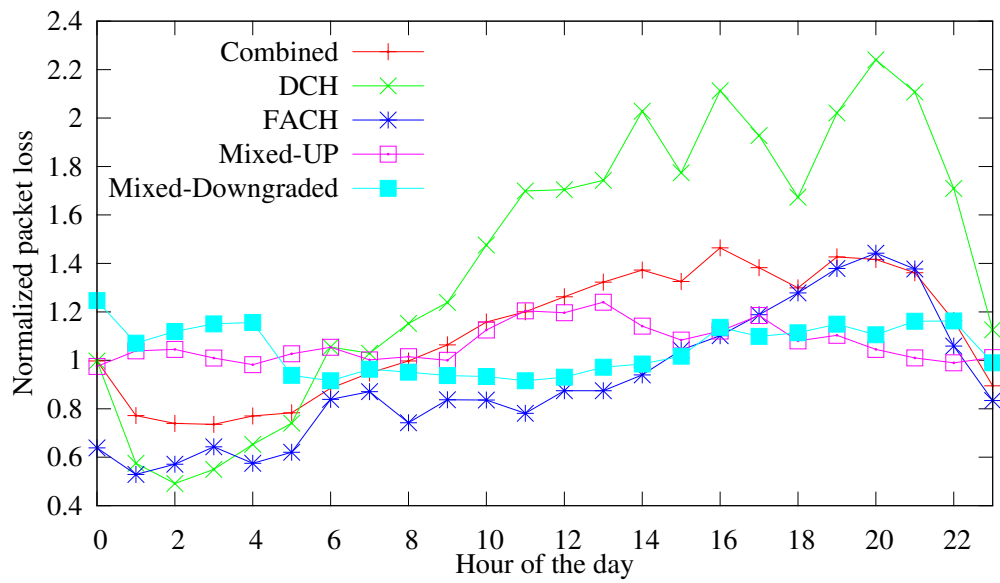
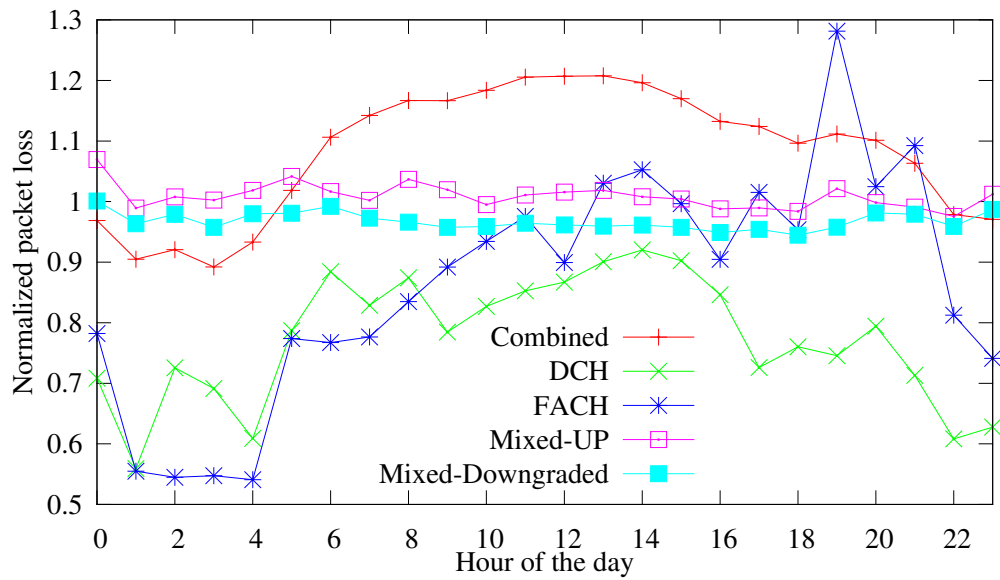


Figure 6.17: Diurnal patterns of loss in Telenor's RAN (top) and Netcom's RAN (bottom). DCH+FACH loss climbs rapidly during peak hours.

egories have similar loss rate regardless of the time of the day. The underlying cause is likely related with forceful RRC state demotions, further described in 6.3.7. In Netcom, the RRC category experiencing most of packet loss, especially during the hours from 14 until 22, is *DCH*. This can be explained by the fact that the majority of our Netcom connections were predominantly on *DCH* during the measurement period due to the administrative traffic sent over them in addition to the UDP ping packets. The same is not true for Telenor, where a significant fraction of our connections were flapping between *FACH* and *DCH* RRC states.

We also note that *Mixed-UP* and *Mixed-Downgraded* categories do not show a strong diurnal component regardless of the network. This suggests that the causes for the observed loss may lie beyond the RAN.

Summary of findings

This section shows that a significant fraction of loss is related to episodes that simultaneously impact a significant fraction of geographically-diverse connections. Further, some loss in *DCH* and *FACH* is characterised by diurnal patterns. More importantly, the identified diurnal patterns seem to be independent of the RNCs the respective MBB connections were controlled by. These findings suggest that congestion in the CN is a plausible cause for loss in *DCH* and *FACH*.

6.3.7 Discussion

In the previous sections, we identified two potential causes of loss: pathological and lossy state transitions; and congestion. Our findings confirm the conventional wisdom that losing packets on the MBB radio interface is rare, especially for reasonably covered stationary connections, thanks to the retransmission-aggressive nature of MBB link-layer protocols. Hence, a MBB network with high loss is indicative of a sub-optimally configured RAN or a congested CN. Congestion can be addressed by re-dimensioning and increasing the capacity of the RAN and CN. However, mitigating loss caused by pathological and regular state transitions requires optimising the RAN configurations. RNCs must not demote users to IDLE as long as they are sending data; such state demotion must strictly be decided by inactivity timers. Also, in flight RRC SDUs need to be buffered during state transitions to account for the mismatch between the SDUs of different channels [8].

Our results inspired Telenor to revisit its RNC configuration to reflect the aforementioned fixes. The top panel in Figure 6.18 shows the CDF of loss percentage

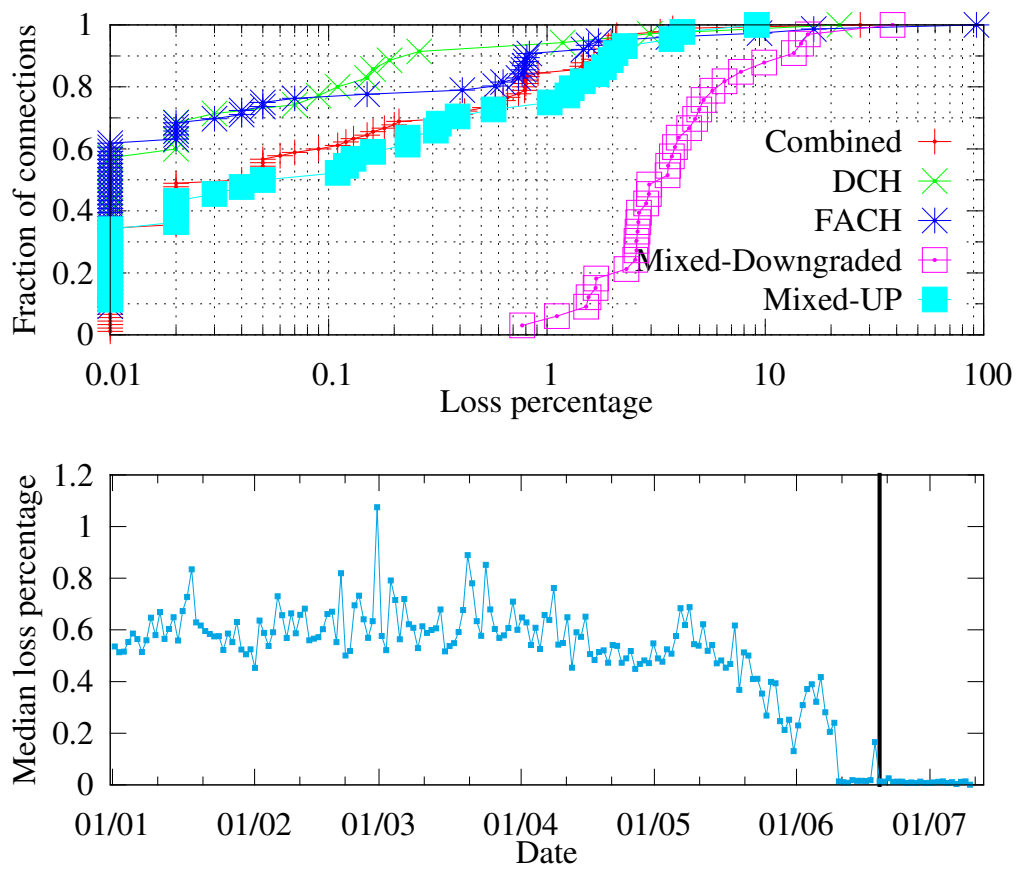


Figure 6.18: Loss rate by RRC state (top) and median loss rate (bottom) for Telenor after the RNC reconfiguration.

after the application of configuration changes in Telenor. The overall loss dropped markedly, and loss in Mixed-UP is similar to loss in FACH and DCH. After the change, the percentage of time Telenor spent on Mixed-Downgraded bins dropped from 19% to only 0.98%. The lower panel shows the median daily loss percent in Telenor measured for all connections, from 01/01/2014 to 10/07/2014. The change was initially applied to only a set of RNCs in the first half of June 2014. The remaining RNCs were patched on 19/06/2014 (marked by the vertical line), where we see a sudden drop in median loss. Following the configuration changes, median loss dropped from about 0.6% to almost zero. This indicates that slight changes to the way the RRC state machine is managed and to the way in-flight data packets are handled during RRC state promotions and demotions can reduce loss significantly. We believe that many operators world-wide would benefit from our observations and the suggested fixes.

6.3.8 Summary of findings

Our results demonstrate that end-to-end measurements can give insights into the performance of cellular network internals. We find that most loss is a direct result of RRC state transitions, both regular and pathological. The remaining loss is mostly due to activity in the CN. Loss exhibits strong diurnal patterns and is related to performance degradation episodes that simultaneously impact a significant fraction of geographically-diverse connections. Our results motivated one of the operators measured, Telenor, to re-examine its network configuration to mitigate loss caused by state transitions. This clearly highlights the importance of independent end-to-end measurement infrastructures like NNE, which allows correlating measurements at different levels and spot potential problems in MBB networks.

6.4 Data plane: investigating excessive delays

One fundamental performance parameter is the RTT between a mobile device and the remote end of its communication. Very large or variable RTTs often signify degraded transport and/or application performance. The behaviour of this delay parameter is affected by a complex interaction of factors including the RAT (2G, 3G, or LTE), access channel type (dedicated vs. shared), channel quality, and traffic level on the RAN and mobile CN.

In this section, we report the results of a measurement study of end-to-end delay in two major Norwegian MBB networks, Telenor and Netcom. We use measurements from the stationary NNE nodes and we quantify and explain observed delay behaviour, including extreme delay events. The delay characteristics for the NNE nodes under mobility is described in Section 7.2.

6.4.1 Basic delay statistics

End-to-end delay in MBB is affected by three factors:

1. RAT: 2G RTTs are $O(100\text{ms})$, 3G and LTE are normally $O(10\text{ms})$
2. wireless standard in use, e.g. HSPA, HSPA+, etc.
3. UE RRC state: (for 3G; LTE only has one state for data transmission). Before transmitting data, the UE attaches to the network and establishes a PDP context with the GGSN. The PDP data structure contains the IP address and other information about the user session. Depending on the traffic pattern, the RNC then allocates a shared forward access channel (FACH = $O(100\text{ms})$) or dedicated radio channel (DCH = $O(10\text{ms})$) for the UE.

This section reports typical RTT values, and their variability, for connections using different RATs and radio resource states. For each connection, we divide the measurement duration into 1-minute bins, classify these bins according to RAT and RRC state (legend in Fig. 6.19). We assign a specific RAT to a 1-minute bin if no RAT change has been reported by the modem. Otherwise, if more than one distinct RAT was observed during a 1-minute bin, we classify the bin as mixed-RATs. In the same way, we assign RRC states to the 3G bins and we do not differentiate LTE RRC states as there is only one non-idle state. We then calculate average and maximum RTTs in each bin. Figure 6.19 shows significant differences in average and maximum RTTs between different RATs and RRC states. Unsurprisingly, LTE and 3G-DCH exhibit the shortest average and maximum RTTs, with maximum 3G-DCH RTTs slightly higher than maximum LTE RTTs. For 2G connections, average RTTs are less than 500 msec in 70% of bins, but they can be as long as several seconds, consistent with the fact that 2G deployments are often in challenging locations where it is not cost-effective to deploy 3G and LTE.

RRC state transitions affect performance, but bins with inter-RAT changes (handovers) have the highest RTTs, also unsurprisingly since handovers can take seconds, leaving packets buffered until the process finishes. Mean RTTs in bins

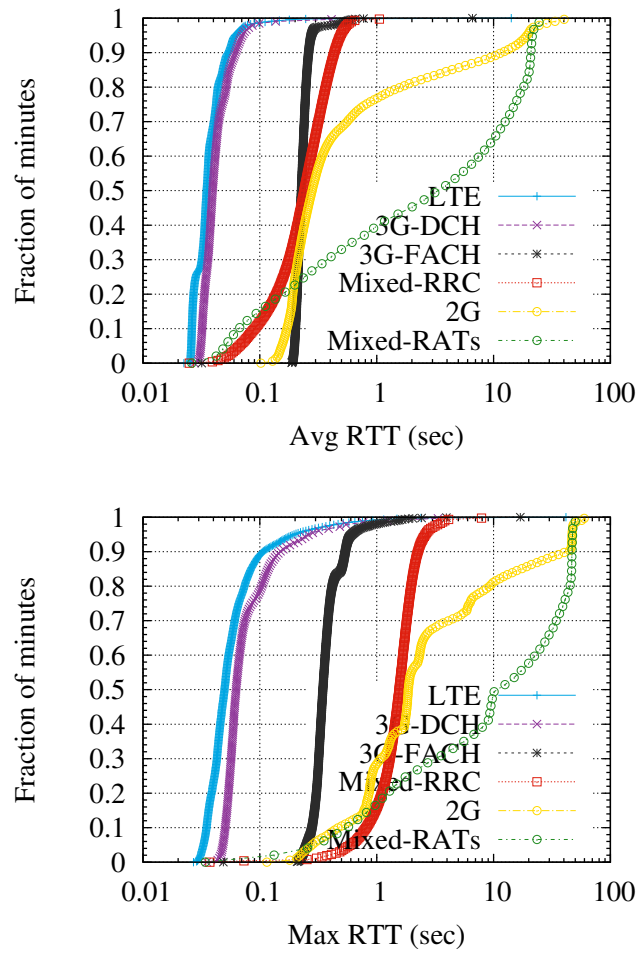


Figure 6.19: Average and max RTTs for stationary connections on Telenor (one-minute bins). Except for bins with mixed-RATs and 2G, RTTs are stable. 3G-DCH and LTE maximum RTTs are double the averages.

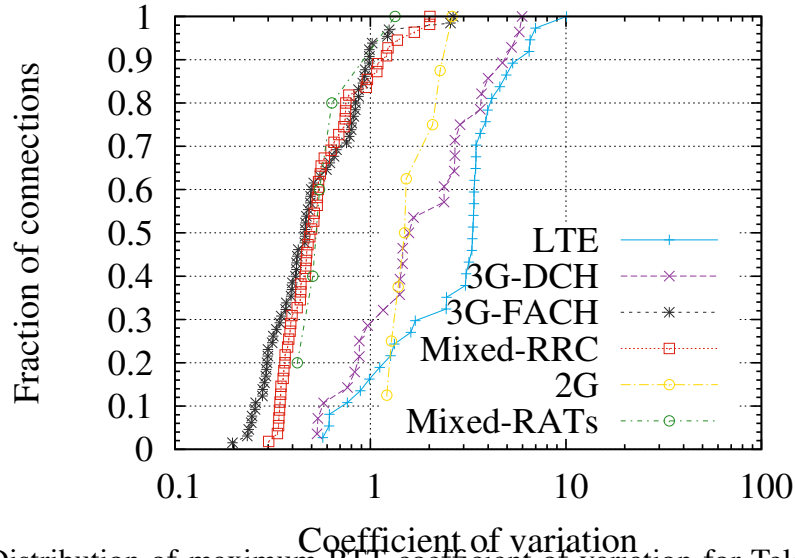


Figure 6.20: Distribution of maximum RTT coefficient of variation for Telenor connections. LTE, 3G-DCH, and 2G connections have highly variable maximum RTTs.

with RRC state transitions are comparable to FACH (shared channel) bins, suggesting that the presence of FACH data in the bin overshadows delay caused by the state transition.

Maximum RTTs in bins with RRC state transitions, on 2G, and with handovers are more than one second, but are clearly outliers; most packets experience far lower RTTs according to the average RTT distributions. Netcom shows similar results, albeit with small quantitative differences.

Variations in RTT

To study whether the high variability in RTTs is common across and within connections, we look at each connection separately, calculating the average and standard deviation of maximum RTTs. We include a connection to a given RAT and/or RRC state sample if the connection contributes at least one day of measurements while on that RAT and/or RRC state.

We use the coefficient of variation, which is the outcome of dividing the standard deviation by the mean, to quantify variability. A coefficient of variation less than one indicates low variability, and above one indicates high variability. Fig. 6.20

shows the CDF of this coefficient for Telenor’s connections for different RAT and/or RRC state combinations. We do not show the same for Netcom connections since the results are largely similar. We measure high variability of maximum RTTs for 2G, 3G-DCH and LTE connections. For instance, the standard deviation is three times the mean for half of Telenor’s LTE connections, which is surprisingly unpredictable for stationary measurement nodes. The variability is lower for connections using 3G-FACH or mixed RRC states, consistent with their max-RTT distributions (Fig. 6.19). Since packets are buffered during RRC state transitions, the maximum RTT in bins with mixed RRC state depends primarily on the length of this transition, which is not highly variable.

FACH (shared access) connections have a much higher baseline, so the magnitude of variability would have to be much higher to yield a large coefficient of variation. For example, if channel quality degradation increases delay by 100msec (i.e. due to link layer retransmissions), delay in a typical FACH connection will increase by $\approx 50\%$ vs. $\approx 200\text{--}300\%$ for 3G-DCH and LTE. Shared (e.g., FACH) connections are more likely to experience congestion delay, but for the two operators we studied we observed much lower packet loss for FACH compared to DCH, which suggests a lack of congestion. One of these operators confirmed privately to us that their FACH channels are over-provisioned. We also found low variations in average RTT, with the exception of 2G, as expected from Fig. 6.19.

In summary, our measurements show that variations in maximum RTTs while on 3G-DCH and LTE are high even at the same location. However, maximum RTTs for 3G-FACH connections exhibit low variability, which might be beneficial for applications that stay mostly on FACH such as M2M applications given that they can tolerate higher delays. We also find that average RTTs measured in a window of one-minute duration are stable and predictable, which shows that most observed episodes of increased delays are too brief to affect statistical averages. We take a closer look at these excessive RTTs in the next section.

6.4.2 Characterising excessive delays

To investigate excessive delays, we identify all 3G-DCH one-minute bins with maximum RTTs of at least two seconds. We compare the maximum RTTs to the RTTs immediately before and after. The preceding RTTs are consistent with typical 3G-DCH RTTs, i.e. $O(10\text{ms})$, but the packet with the highest RTT and the following 2-3 packets appear to arrive at the same time. Figure 6.21 illustrates this RTT behaviour, which results in a triangle pattern when plotted against the

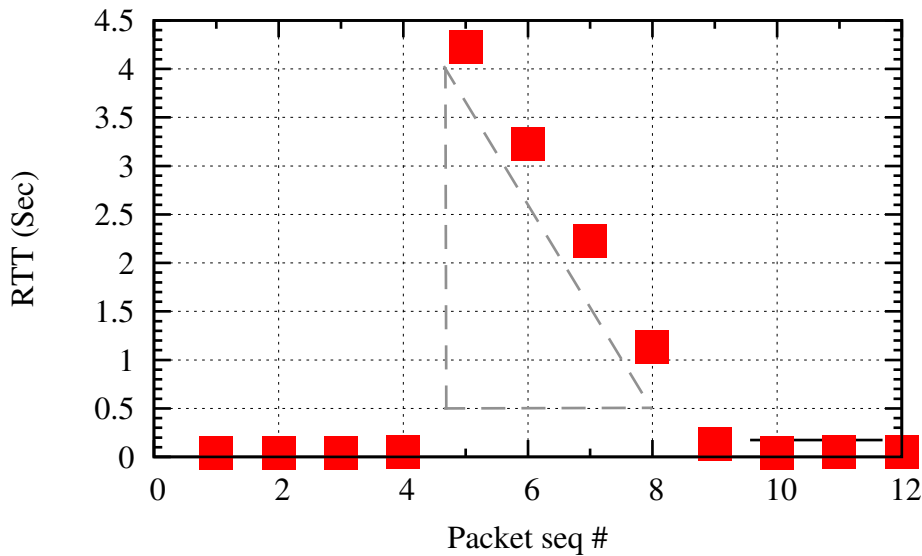


Figure 6.21: Triangle event illustration.

packet sequence number. We call such episodes “triangle events”. The first four packets do not experience excessive delay, but the following four do. The maximum RTT is 4.3 seconds, and the following RTTs monotonically decrease by one second, eventually dropping to the level before the jump. Note that we send our UDP probes every second, which implies that packets five through nine were buffered and then released simultaneously. We next try to thoroughly characterise triangle events.

Detecting triangle events

We say that a triangle event has occurred when we see a sequence of at least two packets, the first of which has an RTT of at least two seconds, followed by a sequence of packets with each RTT decreasing 0.8–1.2 seconds, or only increases or decreases slightly (RTTs within 75% of each other) when the RTTs are large (≥ 1 s). We allow for this large-RTT scenario because we occasionally see such RTTs in the middle of otherwise typical triangle event RTT sequences.

In this analysis we discard events that last longer than 60 seconds, involve packet reordering, or involve more than 60 packet losses, i.e. one minute worth of loss. Our goal is to avoid inflating the number of detected triangle events by including periods that bear similarity to triangle events yet exhibit additional evid-

ence of degraded performance.

We detect triangle events that happened for all Telenor and Netcom connections that contributed at least 10 days of measurements during November 2014. We group events into 5 categories based on RAT and RRC state (when applicable) during the minute in which a triangle event happens: events that occur in bins with RRC state transitions, events in bins with mixed RATs, and events in bins with 2G, 3G-DCH, 3G-FACH and LTE respectively.

Frequency of triangle events

Figure 6.22 plots the CDF of the average number of triangle events per hour for different RAT and RRC state categories for Telenor (top) and Netcom (bottom). We calculate averages by dividing the total number of events in a category by the total time (normalised to hours) spent in that category. The plot shows that triangle events are frequent during RRC state transitions, RAT changes, and 2G, consistent with the associated packet buffering delays (for the first two) and infrastructure challenges (for 2G). A large fraction of 3G-DCH connections, 30% in Netcom and 20% in Telenor, experience more than one event per hour. LTE and FACH appear to suffer fewer triangle events.

Duration of triangle events

Figure 6.23 presents the probability density function of triangle event duration for Telenor and Netcom respectively. We observe distinct modes for most connection classes. For both operators, triangle events affecting 3G-FACH, 3G-DCH, and LTE connections generally last between two and four seconds. The two operators have quite different triangle event duration profiles for connections that experience RRC transitions: tightly clustered around 2.5s for Telenor, and more evenly dispersed around 3s for Netcom. On the other hand, triangle events in bins with RAT changes are longer in Telenor than in Netcom, with modes of 6.5s and 3.2s, respectively. These observed differences between operators confirm that our results are not caused by measurement artefacts related to our hardware. Further, they suggest a room for improvement in network management, i.e., operators can configure their RRC transition and inter-RAT handover times such that they reduce buffering.

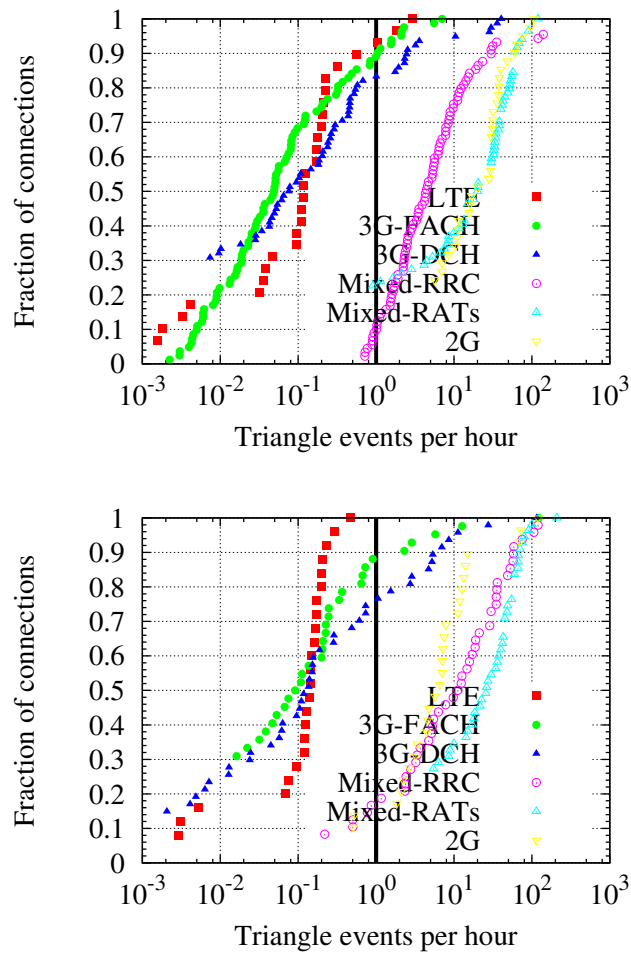


Figure 6.22: The distribution of the frequency of triangle events for Telenor (top) and Netcom (bottom). The solid line marks the fraction of connections that experience more than a single event per hour on average. Approximately one quarter of 3G-DCH connections experience one or more events per hour, on average.

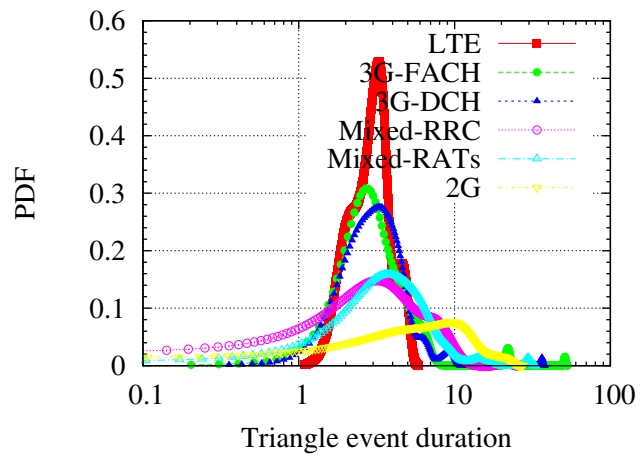
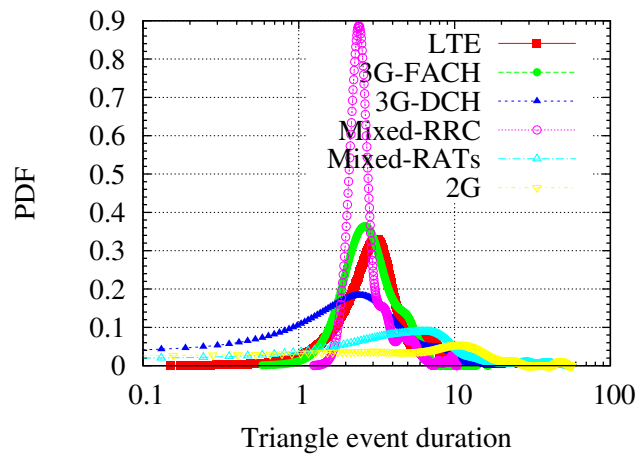


Figure 6.23: PDF of triangle event duration for Telenor (top) and Netcom (bottom).

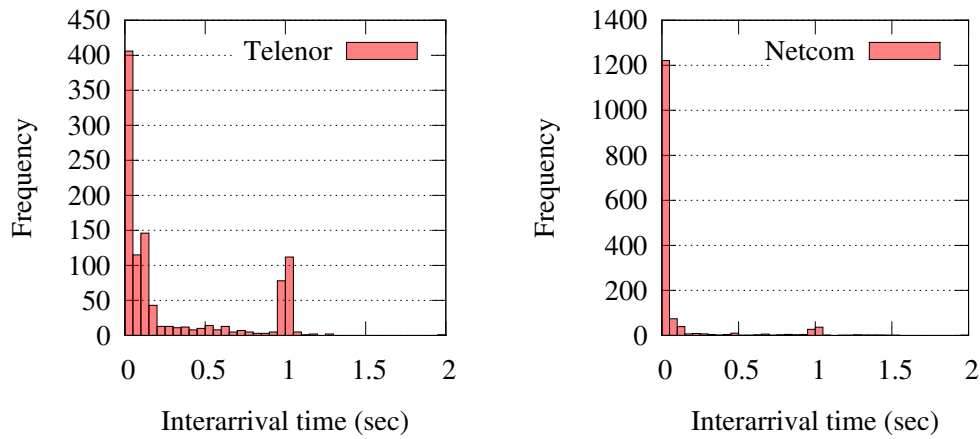


Figure 6.24: Median packet inter-arrival at server during triangle events for Telenor (left) and Netcom (right).

Where does the buffering happen?

To test whether triangle events were more common on the uplink vs the downlink we recorded packet arrival times during triangle events at the server over a three-day period. Figure 6.24 shows histograms of the median packet inter-arrival times during triangle events for connections on 3G. We find a bimodal distribution: zero and one second between packets that make up triangle events arriving at the server, for both operators, with most packets arriving zero seconds apart. Packets arriving essentially simultaneously at the server signifies that they were buffered somewhere on the uplink and released at once. One second spacing between packet arrivals at the server signifies that the packets were not intercepted on their way to the server, as the client sent them one second apart. Thus, triangle events appear to occur both on the uplink and the downlink, though predominantly on the uplink. User equipment transmits packets on the uplink, whereas a cell tower transmits them on the downlink. Cell towers have greater power and more sophisticated coding schemes, and thus can more easily overcome poor signal quality (E_c/I_o and RSSI), resulting in fewer losses, less buffering, and fewer triangle events on the downlink.

Triangle events and signal quality

While some observed triangle events are a direct consequence of the time required to perform RRC state transitions and RAT switches, triangle events also happen at times during which no transitions occur. To gain insight into what causes such events, we correlate the rate of triangle events seen by each connection with its average Received Signal Strength Indicator (RSSI) and signal-to-noise ratio (E_c/I_o or received energy per code bit / interference level). We find an inverse relationship between signal-to-noise ratio and triangle event frequency. A poor signal-to-noise ratio could induce packet loss and link-layer retransmissions. For example, in the data link layer Automatic Repeat reQuest (ARQ) protocol Selective Repeat, a sender must buffer all packets until they are acknowledged. If a series of packets is lost, the sender will maintain a buffer of the unacknowledged packets and release them together, as soon as a packet is acknowledged. Such behaviour matches the one-second spaced decreasing round trip time pattern we observed in triangle events.

6.4.3 Underlying causes

To dig deeper into the underlying causes of excessive buffering episodes, we conduct a set of controlled experiments that monitor the link-layer and radio activity. We focus on buffering in 3G DCH channels not caused by RAT changes or RRC state transitions, thus, we force our modem to stay on 3G and send large (1000-byte) UDP packets, ensuring that we stay on DCH. Finally, since triangle events tend to occur more frequently when signal strength is poor, we place our equipment in an indoor location with $RSSI \leq -103\text{dBm}$ (i.e., the equivalent of one signal bar on most phones). We use Qualcomm eXtensible Diagnostic Monitor Professional (QxDM) [78] diagnostic software to capture link-level messages and events. To test whether buffering happens on the uplink, downlink, or both, we enable logging of received and transmitted UDP packets on our echo server. We ran the experiment for two days to capture buffering episodes in both networks. Our two main observations are:

1. Buffering episodes on the downlink coincide with cell changes. We observe that triangle events occur both during UE handovers to an already known cell following a procedure called HSDPA soft repoint, and during UE handovers to a newly discovered cell [50]. The latter is a two-step procedure that starts with the access network incorporating the new cell into the list of active cells

that the UE can connect with, via an *ACTIVE SET UPDATE* message, after which an HSDPA soft re-point is performed. In both cases, the UE stops receiving RLC protocol acknowledgements for outgoing packets, although these packets are received on the server side without buffering. We observe that in between the cell changes there are no acknowledgements for each sent WCDMA RLC Uplink (UL) Acknowledged Mode (AM) User Plane PDU and also no WCDMA RLC Downlink (DL) PDUs are received. This hints that after the cell change, which involved the change of Primary Scrambling Codes (PSCs), there was no activity on the downlink channel. This lack of acknowledgement indicates that the downlink channel is temporarily unavailable. During this temporary downlink unavailability, incoming frames will be buffered if they are sent in AM, otherwise they will be dropped altogether.

2. Buffering episodes also take place when the UE loses its signalling radio bearer in the midst of the handover. This results in additional negotiation to re-establish the bearer. More specifically, we observe a demotion to WCDMA *LI IDLE* state in the middle of the buffering event and radio bearer reconfiguration of the Enhanced Uplink Channel (EUL). We find that buffering events are the most severe (longest duration) in these cases.

6.4.4 Summary

We have performed a large-scale measurement study of round trip delays in two mobile operators in Norway using static measurement nodes. The results for the moving nodes are provided in Section 7.2. Our work highlights the complex interplay between upper layer protocols and radio and link layer operations.

We saw significant variations in delay both across and within connections, but found that RTTs were mainly a function of the type of MBB technology and data channel used. Correlating these episodes with meta-data captured by our measurement infrastructure indicates that such events are caused by handovers, radio resource state transitions, and heavy retransmissions by physical and link layer protocols. We identified high variability in maximum RTTs measured in one minute bins for 3G-DCH and LTE. We also observe that maximum RTTs can reach several seconds. Investigating these extreme delay episodes, which we term “triangle events,” we find they are caused by excessive buffering on both the uplink and the downlink. Using connection metadata and the QxDM tool, we determine that

these episodes occur due to inter-RAT handovers, RRC state transitions, link-layer retransmissions, inter-cell handovers, and reconfiguration of data channels. We believe that eliminating triangle events, particularly during inter-cell handovers, will be crucial in the near future with the increasing deployment of micro and femto cells.

6.5 Application level reliability

An important aspect of reliability is a stable performance on the application layer. In this section, we look at the stability in performance of two representative applications: HTTP download and VoIP using the Session Initiation Protocol (SIP).

6.5.1 HTTP download test

Much of the traffic in MBB networks goes over the HTTP protocol, which is used for both web browsing and video streaming. Here, we report results from a simple experiment where a 1 megabyte file is downloaded from a server using the HTTP GET method. For each test iteration, we record the time it took to start the data transfer, the total time to complete the download (if successful), throughput, and any error code. For each connection, the test was repeated once per hour for a duration of 3 weeks, giving a total of up to 504 runs per connection. Based on these measurements, we report on two different metrics: the probability of successfully completing the transfer, and the probability of achieving a download rate of at least 1 Mbps.

Note that our aim here is to look at the stability of the application layer performance, and not to measure the maximum achievable throughput in the connection, which would require a different approach.

Figure 6.25 shows the fraction of download attempts that could not be successfully completed. We observe that the fraction of failing attempts is significantly higher in Telenor than in the other networks. 55% of Telenor connections experience a failure rate higher than 5%, and 12% experience a failure rate higher than 10%. Ice, on the other hand, sees very few unsuccessful downloads, with only 9% of connections having a failure rate above 1%.

Looking at the error codes for the unsuccessful transfers, we find that the dominating reason for a failed download in all networks except Telenor is the failure to establish the TCP connection. Correlating erroneous transfers with the data from our connectivity test, we observed that these unsuccessful TCP handshakes

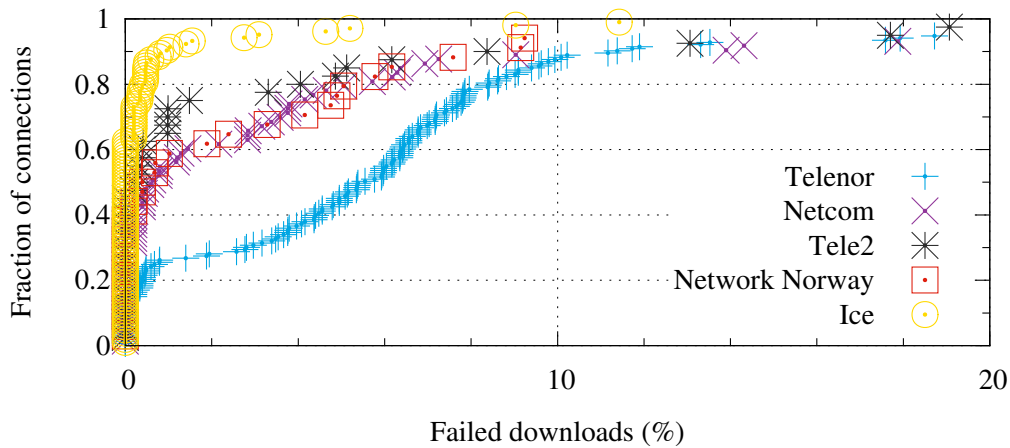


Figure 6.25: Failure rates in HTTP download tests.

happen during times with high packet loss. For Telenor, however, the connection broke *after* the initial TCP handshake in 74% of the failing attempts. Such errors are largely responsible for the difference between Telenor and the other networks. Looking closer at these events, we find that they happen when the connection can not be promoted from the CELL_FACH to the CELL_DCH state. The download traffic will normally trigger such a promotion, but when this does not happen, the modem responds by resetting the connection, and the download fails. There is a clear diurnal pattern in this type of failures, which makes it natural to believe that they are related to congestion in the FACH channel and/or insufficient resources for a promotion to CELL_DCH.

Figure 6.26 shows, for each connection, the fraction of runs where the achieved throughput was less than 1 Mbps. Only 3G connections are included in this plot. We observe that in all UMTS operators, most connections achieve this download rate most of the time. In Netcom, Tele2 and Network Norway, around 90% of connections achieve at least 1 Mbps 90% of the time or more. Telenor achieves this rate less often. We believe that this is caused by the higher loss rate observed in Telenor and how it interacts with TCP congestion control. Ice achieves far lower throughput than the UMTS operators. Around 20% of the connections never achieve the target download speed, while only 19% of connections meets the target in 90% of the cases. From Figures 6.25 and 6.26, we conclude that Ice offers a slower but more stable download performance than the UMTS networks.

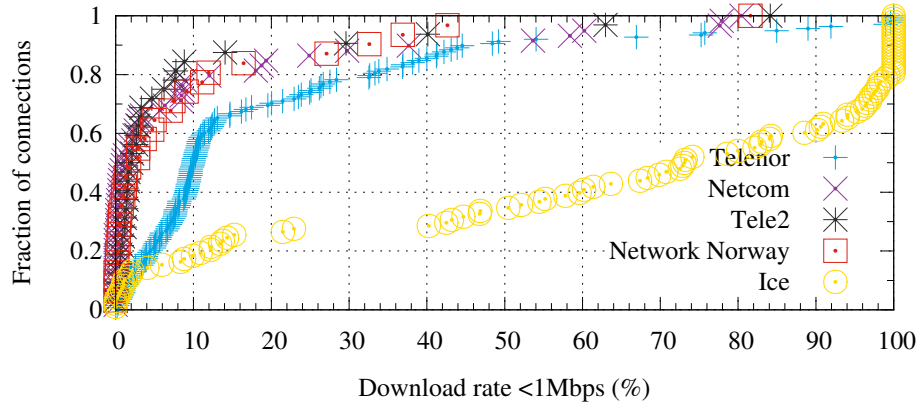


Figure 6.26: Probability of achieving less than 1Mbps HTTP throughput.

6.5.2 VoIP test

Initiating VoIP calls over MBB connections is becoming increasingly popular, thanks to the rise of applications such as Viber [98] and WhatsApp [47]. The ability to initiate and complete VoIP calls are therefore important aspects of the user-experienced reliability. To assess this ability we design a simple experiment that faithfully emulates a real VoIP call. Our test consists of a custom-made client that runs on the measurement node and an Asterisk Private Branch Exchange (PBX) [27] hosted at the NNE backend. The client initiates a VoIP call to the PBX using SIP, then it uses Real-time Transport Protocol (RTP) to play a one minute long audio file. Upon completion the PBX replays the same file back to the client and terminates the call. The audio file is encoded as G.711 [1] resulting in a sending rate of about 50 packets per second. For each connection, we run this experiment once per hour for one week, giving us 168 calls per connection.

Figure 6.27 illustrates the fraction of calls that could not be successfully completed. We observe that the failure rate is significantly higher for Telenor. 30% of Telenor connections have a failure rate higher than 15% compared to only 4% of Tele2 connections. 34% and 21% of the failures for Netcom and Tele2 respectively happened during the call initiation phase (i.e. the SIP INVITE fails). For Telenor and Network Norway, however, this percentage drops to 13% and 14.9% respectively. The remaining failures happened after the call started successfully. We believe that the explanation for these dropped calls is the same as for the unsuccessful HTTP downloads discussed above.

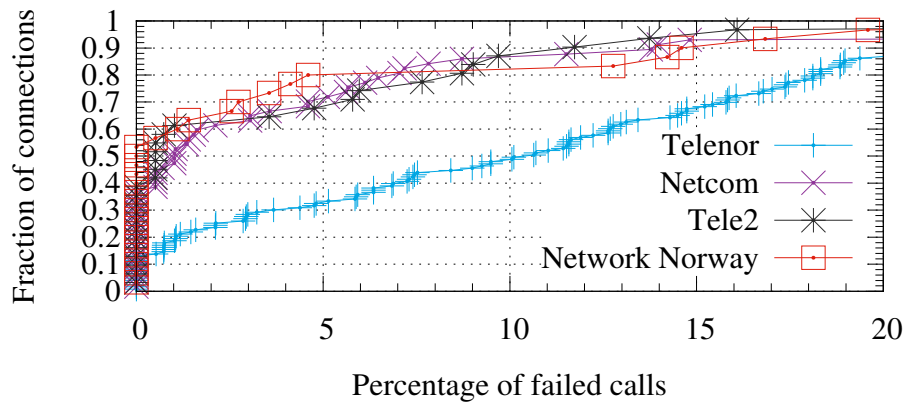


Figure 6.27: Failure rates in VoIP tests.

6.5.3 Summary of findings

This section has shown that short pauses in packet forwarding that are likely caused by the lack of available radio resources can lead to significant problems for applications, such as dropped calls and failed downloads.

6.6 Robustness by multi-homing

So far, our discussion has focused on the experienced reliability in each MBB network separately. A related question is how much reliability can be increased if an end device can be simultaneously connected to more than one MBB network. Such multi-homed devices are becoming increasingly popular, for example as mobile gateways that provide WiFi service on public transport.

The potential for increased reliability through end device multi-homing depends heavily on whether coverage and failure patterns are largely independent across operators. If connections to different operators tend to fail simultaneously due to the same underlying reasons, there is little to be gained from a redundant connection. Our discussion in this section focuses on three important aspects of cross-correlation between operators: coverage, connection failures and loss events.

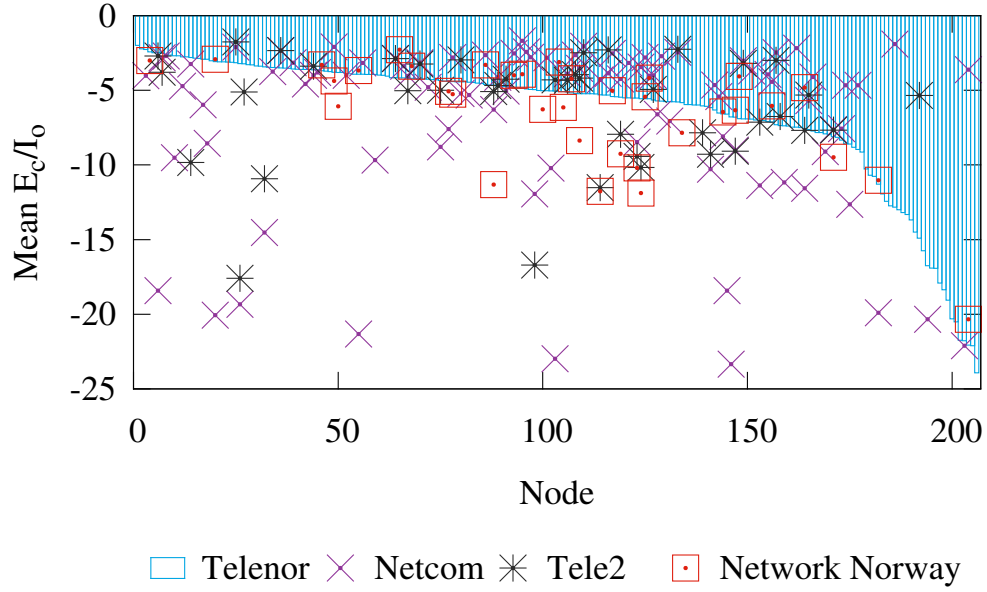


Figure 6.28: Average E_c/I_o values for connections in different operators.

6.6.1 Correlations in coverage

We first look at whether different operators are similar with respect to signal quality at a given location. NNE measurement nodes are distributed widely across Norway, and we believe they give a reasonable representation of the indoor signal quality that can be expected in each operator.

Figure 6.28 gives an overview of the E_c/I_o values for all UMTS connections in all nodes in our dataset. The values shown are averages across the whole measurement period (for most connections the variation is small). Since almost all nodes have a Telenor connection, we sort the nodes based on the E_c/I_o of the Telenor connection, and plot E_c/I_o values for the other connections where available.

We observe that there does not seem to be a strong correlation between operators in general. As expected, we find that for some nodes, Netcom/Tele2 and Telenor/Network Norway have pairwise very similar E_c/I_o due to their national roaming arrangements. Calculating the Pearson correlation coefficient for each pair of operators confirms the visual impression from Figure 6.28. The correlation coefficient for the pairs of operators that do not share the same RAN is between -0.10 (Netcom and Network Norway) and 0.25 (Telenor and Netcom). The correlation is evidently higher when the respective pair of operators share the same

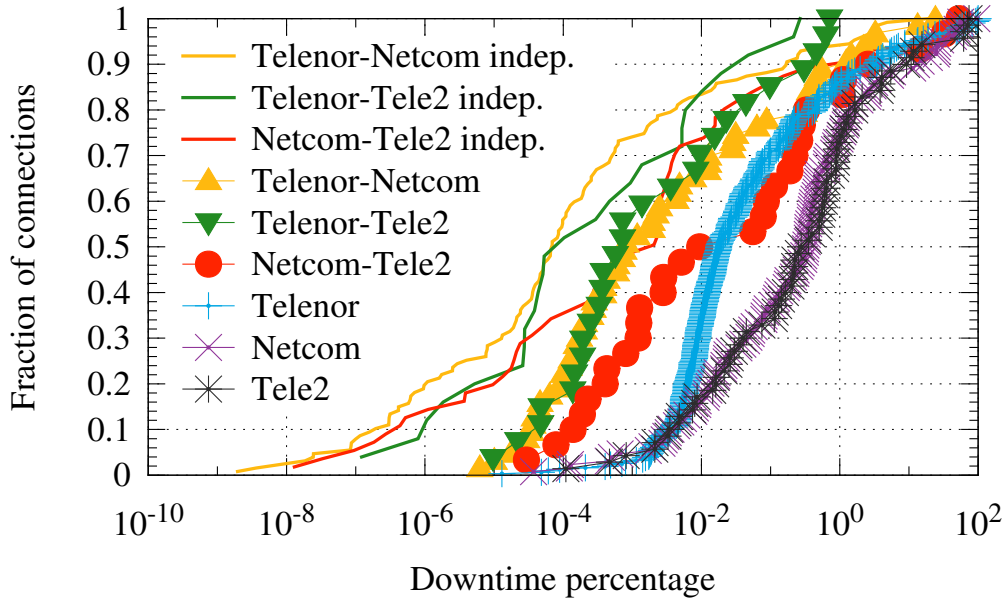


Figure 6.29: Common downtime across operators.

RAN, that is 0.47 for Tele2 and Netcom and 0.75 for Telenor and Network Norway. We also performed a similar analysis for RSSI values, which gave similar results. These findings are positive from a robustness perspective, since they indicate that there is a significant potential gain from provider multi-homing.

6.6.2 Correlations in downtime

Next, we revisit the connection failures discussed in Section 6.1, to see how downtime could be reduced if the end user is able to connect to more than one operator. For each pair of connections from the same node, we identify the time periods when both connections were unavailable (i.e., they had no PDP context). The resulting downtime represents a lower limit on the downtime that can be achieved if the end system is able to exploit both connections in parallel.

Figure 6.29 shows the downtime for three operators (repeated from Fig. 6.1), and the combined downtime for all pairs of connections from the same operators. Other operators and pairs are excluded from the graph for readability. For comparison, we also plot the expected combined downtime for connection pairs under the (false) assumption that they fail independently. This value is calculated by multiplying the downtime rates of each connection, and serves as a lower limit on

the downtime we can expect by combining operators.

We make two main observations. First, there is a large potential for reducing downtime through multi-homing. 60% of the nodes achieve 99.999% uptime when combining Telenor and Tele2 connections; and 55% of the nodes achieve the same when combining Netcom and Telenor connections. These reductions in downtime are arguably surprisingly high, given that connections from different operators often share many potential sources of failures, such as local weather conditions, cell towers, power, or even transmission. We see that the measured combined downtime is not very different from the theoretical downtime assuming independence. Second, the reduction depends heavily on selecting the right pair of connections. As expected, operators that often camp on the same RAN show a much larger correlation in their failure pattern, and the gain from combining such connections is limited. As shown in Fig. 6.29, there is often little or no gain in combining connections from Netcom and Tele2.

6.6.3 Correlations in loss

Finally, we look at correlations in loss between networks. We base our analysis on the same 5 minute intervals used in Section 6.2. Let $P(A)$ denote the (empirical) probability that a connection A has a loss rate higher than 10% in a given 5 minute interval, and $P(B)$ be the same for a connection B . We calculate the conditional probability $P(A|B)$ for each pair of connections from the same node, and compare it to the unconditional probability $P(A)$. If the conditional probability ratio $R = P(A|B)/P(A)$ is close to 1, it means that connection A and B fails largely independent, while a high R means that they tend to fail together. Note that by Bayes' law, $P(A|B)/P(A) = P(B|A)/P(B)$.

Figure 6.30 shows R for all pairs of connections at the same node, grouped by operators. We require that A and B have at least one full week of overlapping measurements to be included in the analysis. Note that the number of events where both connections experience high loss may in some cases be very low, so the graph should be interpreted with some care. Some observations are still clear. First, connections from different networks are not completely independent. In between 25 and 70% of connection pairs, the likelihood of high packet loss in connection A more than doubles when connection B experiences high loss. In between 8 and 35% of the cases, the likelihood increases more than 10 times. There are, however, clear differences between operator pairs. Not surprisingly, the strongest pairwise dependence is for Netcom/Tele2 and Telenor/Network Norway. The weakest de-

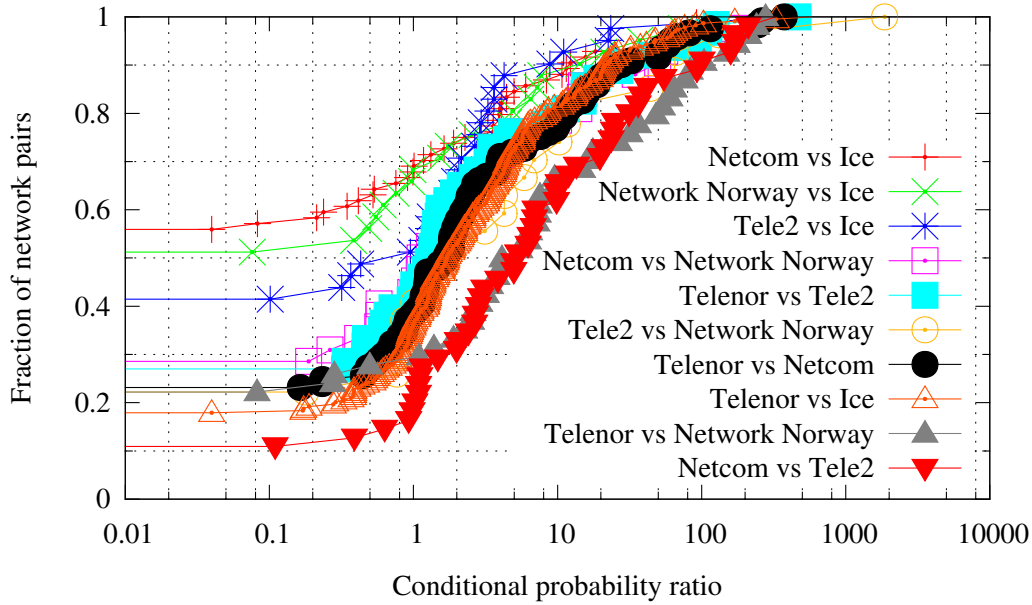


Figure 6.30: Conditional probability of loss events.

pendence is between Ice and the other operators. This might also be expected, since Ice operates with a different technology, a much lower frequency band, and has a different customer mix (and therefore traffic pattern) than the other operators.

In Section 6.3 we showed that apart from RRC state transitions, the remaining loss is mostly due to activity in the CN. Figure 6.30, however, still shows relatively strong correlation for connection pairs from the operators that do not share CN or RAN. One of the reasons could be the fact that in Section 6.3 we used loss threshold of 3%, whereas here we use 10%. To this end, the higher conditional probability ratio values shown in Figure 6.30 could also be coming from the common lossy episodes that happened either beyond the CN or in the shared transmission lines between the base station and the CN.

6.6.4 Summary of findings.

The results in this section indicate that there is a large potential for increased reliability through multi-homing to several MBB networks. There are generally clear differences between operators in signal quality at a given location, and exploiting two connections in parallel can potentially give 99.999% availability in many cases.

6.7 Summary

In this chapter we have presented a large-scale study of reliability in 5 Norwegian MBB networks for the stationary connections. We used the framework proposed in Chapter 3 and focused on a few selected metrics at each level in the framework. This study was performed on the NNE infrastructure, with dedicated measurement nodes.

The measurements presented here have demonstrated that there are clear differences in reliability between operators, and that these can be identified and characterised by end-to-end measurements. Networks vary in the stability of connections, in packet loss and delay patterns, and in their ability to support popular applications.

We find that most loss is a direct result of RRC state transitions, both regular and pathological, while the remaining loss is mostly due to activity in the CN. Loss exhibits strong diurnal patterns and is related to performance degradation episodes that simultaneously impact a significant fraction of geographically-diverse connections. Our study of excessive delays highlighted the complex interplay between physical and transport layer protocols, and showed that the majority of excessive delay events are caused by buffering on both the uplink and the downlink.

We have shown how end-to-end measurements can give insights into the performance of cellular network internals and be used to identify failures and performance problems that are not necessarily captured by the operators' monitoring systems. Our results motivated one of the operators measured to re-examine its network configuration to mitigate loss caused by state transitions. We further showed that using two MBB connections from distinct operators in parallel can potentially give 99.999% availability. This clearly highlights the importance of independent end-to-end measurement infrastructures like NNE, which allows correlating measurements at different levels and spot potential problems in MBB networks.

Chapter 7

Reliability and performance: mobile scenario

7.1 Classification of loss under mobility

The ability to deliver data packets as reliably as possible is arguably one of the most important quality metrics in MBB networks. Excessive and even sporadic packet loss worsens the user experience significantly. It degrades the performance of reliable transport protocols, increases retransmissions, and ultimately degrades application performance.

Assessing and mitigating packet loss is an important step for improving MBB performance. There are, however, many potential causes of loss, which makes characterising and understanding loss a non-trivial task. This task becomes particularly daunting under mobility. Moving connections experience varying signal quality, cellular handovers, potential changes in radio technology, just to name a few.

In this study, we perform a longitudinal practical investigation of packet loss under mobility in operational MBB networks using end-to-end measurements. We look at the loss characteristics in mobile networks when UEs are physically moving. We dissect and seek to explain why loss is much higher under mobility.

To this end, we use over half a year's worth of measurements from six measurement nodes that are placed on board regional and inter-city trains in Norway. We aggregate the measurement data into 5-minute bins in the same way as we did for the stationary scenario. We use this data to compare and characterise loss under mobility to stationary scenarios. We further leverage connection state in-

formation to identify the underlying causes of loss. In most of our analysis, we do not differentiate between movement at different speeds, since we observe that this has a limited effect on packet loss (see 7.1.7). Only less than 4% of our 5-minute bins have the average speed of 100 km/h or more, whereas it has been shown that the effect of speed alone on packet loss is negligible for train speeds below 150 km/h [62].

Our measurements and analysis give insights into the characteristics and causes of packet loss under mobility. In summary, this study makes the following contributions:

1. We present the most comprehensive study of loss in MBB networks under mobility. Using over half a year's worth of measurements and data points from diverse geographic locations, we are able to pinpoint causes of loss under mobility and derive a classification methodology.
2. We demonstrate that performing end-to-end active measurements in conjunction with collecting connections' metadata can help dissecting the most complex MBB scenarios.
3. Our results single out technology handovers and coverage holes as the main causes of loss under mobility. We use these insights and other findings to identify potential areas for improvement.

7.1.1 Measurement setup

The measurement setup used in this study is the mobile subset of the NNE and consists of six measurement nodes placed on six regional and inter-city trains operating in Norway. Figure 7.1 shows the routes covered by these trains, which includes a reasonable mix of urban and rural areas. In this study we limit ourselves to studying two UMTS operators, Telenor and Netcom, because these operators were the only operators that provided LTE service and ran their own radio access and core networks during the period of the study.

We aggregate the data into 5-minute bins and calculate loss percentage for each bin and use the GPS location data from the train's fleet management system to identify the location of NNE measurement nodes and trains speed during the measurements. The GPS locations are updated every 10 to 15 seconds in the fleet management system.



Figure 7.1: Map of the train routes.

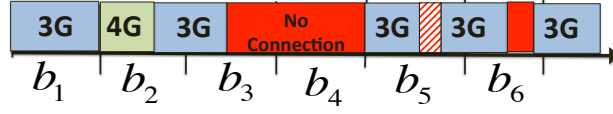


Figure 7.2: Typical sequence of connectivity and coverage conditions that MBB connections experience as they move.

7.1.2 Measurement scenario

Figure 7.2 shows a typical sequence of connectivity and coverage conditions that MBB connections experience as they move. RAT can be constant during a whole bin or several consecutive bins as in b_1 . A bin may involve several horizontal handovers, that is changes of the serving cell. Some bins involve inter-RAT handovers (e.g., handover from 4G to 3G in b_2). All types of handovers are well defined procedures that should normally last a couple of seconds and degrade the user experience negligibly. Connections may suffer from lack of coverage, which leads to a complete loss of connectivity for several minutes (e.g., no connection period extends from the mid of b_3 to the end of b_4). In these periods, a connection loses its PDP context (3G) or EPS bearer (LTE). Consequently the connection loses its IP address. When a connection breaks, software on the nodes immediately checks if there is coverage and tries to reconnect. Otherwise it waits until coverage becomes available. Connections may also experience temporary loss of connectivity that is immediately rectifiable e.g., the short disconnection during b_6 . These episodes can be caused by temporary lack of coverage (i.e., coverage holes) or due to the interplay between mobility patterns and handover procedure decision and duration. For example, the modem may start a handover to a new cell, a procedure that involves current and neighbour cells, but it loses connectivity to the current cell before completing the handover. Another cause can be failures during inter-RAT handovers, which are known to happen [62]. Finally, connections also experience periods with a brief lack of service that are immediately followed by a service restoration without inter-RAT handover or context reset like the shaded area during b_5 . Note that during periods with lack of service, connections typically remain attached to the network and appear to have a PDP context (EPS bearer).

In this section, we are interested in measuring users experience as nodes move and have connectivity. Accordingly, we divide the measurement bins into two groups. 1) bins where users experience lack of coverage, and 2) bins where users have coverage but may experience a brief lack of connectivity that is restored by

an immediate reconnection attempt. All bins in the first category (58% and 54% of bins in Telenor and Netcom, respectively) are discarded in the remainder of this section.

7.1.3 Data curation

UE and connection managers always try to cope with the varying coverage and connectivity conditions by quickly detecting lack of connectivity, attempt to reconnect, or reset the wireless device altogether. The interplay between varying signal conditions and UE hardware is non-trivial and it may sometimes render the connection unusable. We believe that some failure situations are caused by specific measurement and system artefacts; a different system or hardware may cope better or worse. Next we describe these artefacts in more details.

Sometimes modems become unresponsive and are eventually ejected by the operating system, resulting in a disconnect and probably packet loss before the ejection. In some other cases, the PDP context (EPS bearer) might seem to be operational, but IP packets cannot be sent or received until the connection is re-established. We refer to these connections as stale connections. We verify that connections become stale when the network attempts to reset long-lived PDP contexts (EPS bearers), and it fails half-way through the process without actually resetting the context (bearer). As a result, the operator's firewall drops all incoming packets from these connections, causing 100% packet loss during these periods. This clearly illustrates an interplay between the configuration parameters and management procedures imposed by individual network operators, and UEs used by NNE.

Other artefacts include server-side failures and measurements with misreported metadata. For example, the modem reports that it is on LTE while at the same time reporting 3G-specific metadata such as E_c/I_o or Received Signal Code Power (RSCP). This typically happens when the modem delays sending metadata because it is busy with processing control traffic.

To be able to cope with the aforementioned anomalies, we impose a number of filters on the dataset. This leaves us with only measurement data that is supported by clean metadata. Next, we describe our filters:

1. We remove all 5-minute bins that involve one or more connection reset such that the connection did not recover after a single retry.

2. We remove all 5-minute bins with 100% packet loss to avoid stale connections and cases where the modem is stuck and yet appears operational to the OS. By doing that, we risk excluding some legitimate loss events that are caused by equipment failures and maintenance activity. These events are, however, outside the scope of this study since we are interested in what users experience on a daily basis and not rare or scheduled events.
3. We look only at bins where the train was moving, and we require at least one available GPS reading in a 5-minute bin in order to determine this. We impose the average speed of the available readings to be greater than zero. To check for the cases when the train was predominately still during a 5-minute period, we imposed larger average speed thresholds and observed similar results.
4. We remove all 5-minute bins that coincide with known server-side maintenance.
5. We keep only 5-minute bins where we have metadata reports for at least 4 of the 5 minutes. These reports are acquired by polling the modem at the beginning of each minute.
6. We keep only bins with known RATs and valid combination of RAT and RAT-specific metadata.

After curating the initial data set, we have 63837 5-minute bins, 38417 from Telenor and 25420 from Netcom.

7.1.4 Loss in different RATs

The measurement nodes used in this study will always try to connect to the highest available RAT. That is, they will prefer LTE over 3G over 2G. To investigate loss in different RATs, 5-minute bins are divided into 3G, LTE and mixed. Mixed bins are bins where the node was connected to more than a single RAT. We do not include bins spent fully on 2G in our analysis, since there are relatively few such bins, and both loss rates and connection stability are much worse in this RAT. 2G bins are experienced mostly in challenging areas (with limited coverage), and therefore our measurements are not representative for normal 2G behaviour. In particular, our dataset contains only 190 and 148 2G bins for Telenor and Netcom, respectively.

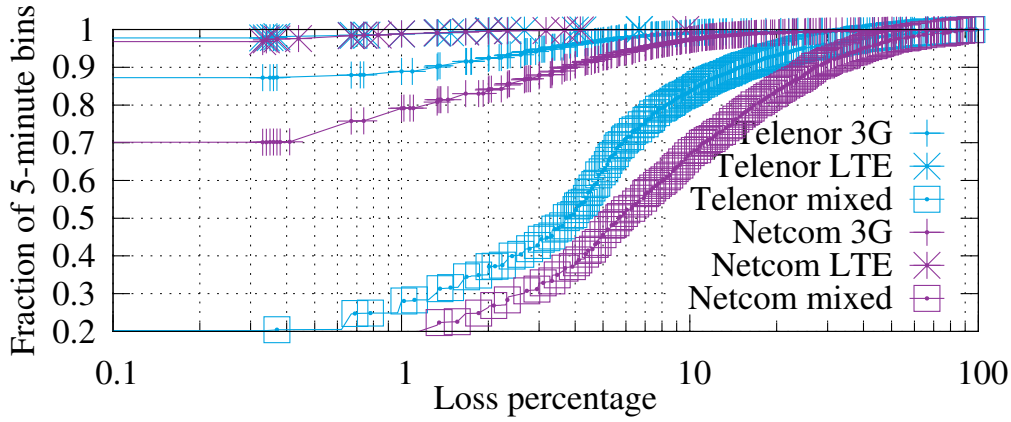


Figure 7.3: Loss rate for 3G, LTE and when a RAT change is involved (mixed). Most loss happens in mixed category, LTE performs the best.

In these bins, the average loss rate is between 17% and 18.4%, while the median loss rate varies from 11.6% to 13%.

Figure 7.3 shows loss for different RATs when the measurement nodes are moving. We first observe that loss rate is higher in 3G than in LTE. Less than 4 % of LTE-only bins experience loss, while 14 % (Telenor) to 40 % (Netcom) of 3G-only bins experience loss.

The most striking observation from Fig. 7.3 is, however, the much higher loss rate in mixed bins. 86% (Telenor) to 94% (Netcom) of bins with RAT changes also have packet loss. In 16 % (Telenor) to 33% (Netcom) of bins, the packet loss is over 10%. This indicates that inter-RAT handovers is a major source of loss in MBB networks. We perform an in-depth analysis of this loss in Section 7.1.6.

7.1.5 Classifying loss under mobility

In order to structure our investigation of loss, we start by classifying all 5-minute bins according to the state of the connection in that bin. We perform this classification in a hierarchical fashion, as shown in Figure 7.4. This classification captures the connectivity states shown in Figure 7.2 and isolates independent conditions, thus reducing the complexity of identifying potential causes of loss.

The root of the tree contains all bins where the measurement nodes have radio coverage as discussed in Section 7.1.2. Inspired by the observation in Figure 7.3 that loss is much higher when there is a RAT change, we first split the bins into

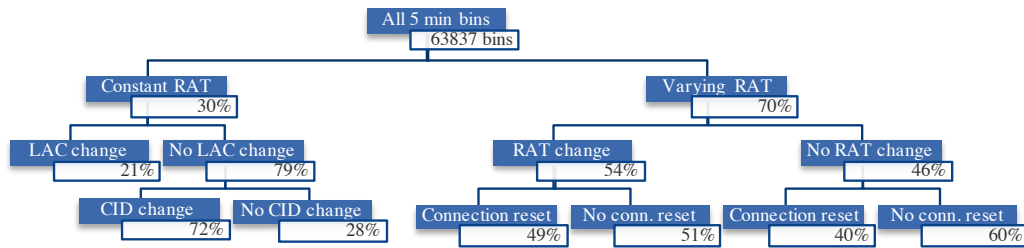


Figure 7.4: Classification of loss. The numbers given are percentages of lost packets relative to the parent category. In total, there are 63837 5-minute bins in the dataset. About 229992 packets were lost during these bins.

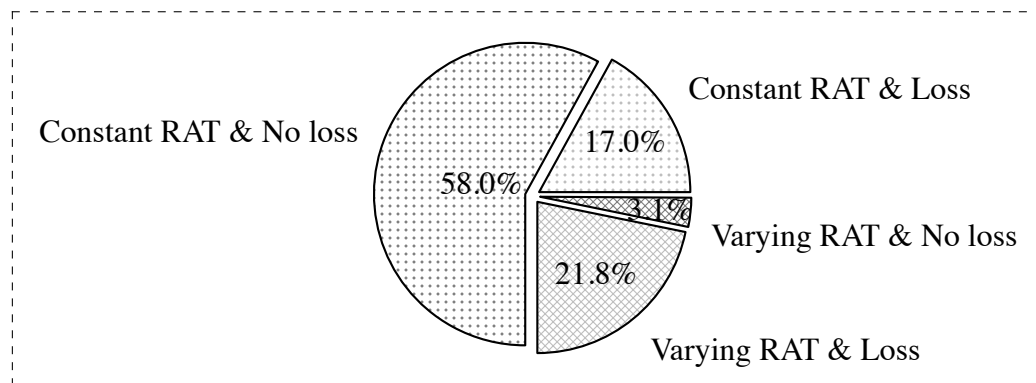


Figure 7.5: The percentages of lossy and non-lossy 5-minute bins for constant and varying RATs.

constant and varying RAT. Constant RAT bins are characterised by a single RAT throughout their duration. Varying RAT bins, however, involve more than one RAT or a short lack of service. 30% of all loss occurred during bins with a constant RAT, while in 70% of loss occurred during bins where the RAT changed at least once. Loss rate is on average seven times higher in bins with varying RAT.

Figure 7.5 shows the distribution of lossy and non-lossy bins for constant and varying RAT cases. The percentage values shown are relative to all bins (they add up to 100%). We observe that a clear majority (more than 3/4) of bins with constant RAT experience no loss. On the other hand, almost all bins with varying RAT involve packet loss.

Bins with varying RAT are further divided into bins where the connection is attached to at least two RATs, and those where we only observe one RAT. Note

that a bin might still be classified as varying even if we only observe one RAT: this means that the modem reported no available RAT at least once during the bin. This behaviour is a normal part of a RAT transition, but it also sometimes appears without a resulting RAT change.

Recall from Section 7.1.2 that we only include 5-minute bins without connection resets, or where a single reconnection attempt immediately restores connectivity. The varying bins with and without RAT changes are further subdivided according to whether there is a connection reset in the bin. We separate bins with connection resets because we believe that they are characteristically different from the rest of the varying bins. These resets are likely caused by small coverage holes and failures of the handover procedure. We observe from Figure 7.4 that all varying RAT categories are responsible for a significant share of the overall loss. Loss in bins with varying RAT is further explored in Section 7.1.6.

The constant RAT bins are further classified according to whether there is a LAC change in the bin, and if not so, whether there is a CID change. The intuition behind this classification is a hypothesis that horizontal handovers (change of LAC or CID) is an important source of loss in constant RAT periods. The numbers in Figure 7.4 confirms that this is the case: 78% of loss in constant RAT periods happen in bins with either LAC or CID changes. Section 7.1.7 provides a more detailed investigation of causes of loss in constant RAT periods.

7.1.6 Varying RAT

In this section, we analyse loss in bins with varying RAT, which constitutes 25% of all bins and are responsible for 70% of the overall loss. As mentioned in Section 7.1.5, bins with varying RAT come in two forms. First, bins with one or more inter-RAT handovers; that is, we observe more than a single RAT in the bin. Second, bins with no handovers but a glitch in the service. The second class of bins are characterised by one predominant RAT (2G, 3G or LTE), but with the presence of one or more *No service* episodes throughout the bin.

Figure 7.6 divides bins with varying RAT according to the presence of RAT changes and loss. A large majority of bins with inter-RAT handovers, about 92%, involve packet loss. This fraction is slightly smaller, about 80%, for bins with no RAT change. In the following subsections, we will look at both scenarios and investigate possible causes of loss.

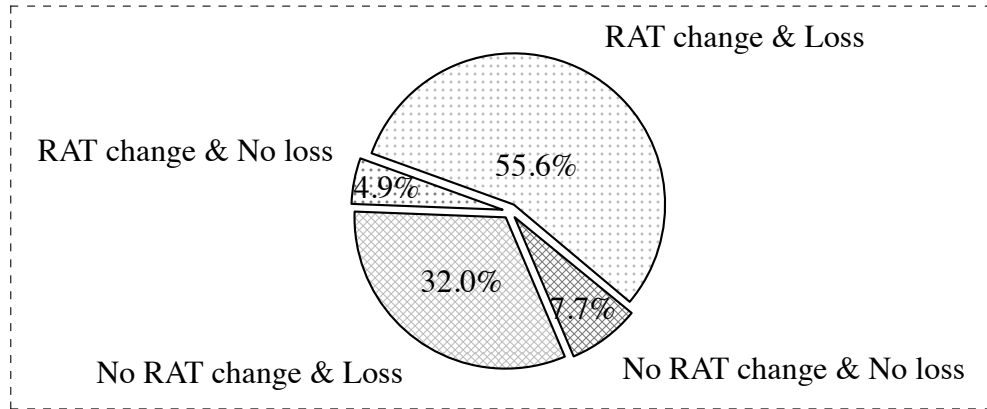


Figure 7.6: The percentages of lossy and non-lossy 5-minute bins with varying RATs split by whether more than one RAT is observed.

Loss during periods with RAT changes

The modems in the setup are configured to automatically select the highest available RAT. In large cities, LTE is almost always available, while outside the metro areas, 3G is the predominant RAT. Inter-city trains cross rural areas where conditions can vary from strong 3G to weak 2G signal to no coverage. RAT changes, or inter-RAT handovers are based on the UE neighbor cell measurement reports, which are regularly sent to the network. Based on these reports, the network can initiate a handover from one RAT to the other. This typically happens when measurements show that the signal and the interference levels from the current cell in RAT A are worse compared to levels of RAT B. It can also happen when the UE moves into the range of a new cell or cell sector that supports RATs different from the current RAT. The handover process involves multiple steps involving the UE, the RAN and in the CN. These steps vary depending on the type of handover, e.g., from 3G to 2G, from LTE to 3G, etc. Further, inter-RAT handovers are not always seamless. For example, sometimes we observe RAT changes that lead to a connection reset. Moreover, it is quite common to lose several packets right before the connection breaks.

Figure 7.7 shows the split of bins containing an inter-RAT handover, according to the presence of loss and connection resets. The main observation is that bins with inter-RAT handovers involve packet loss independent of whether there is a connection reset or not. Almost all (99%) bins with RAT change and connection resets include packet loss compared to 88% of the bins with RAT changes but

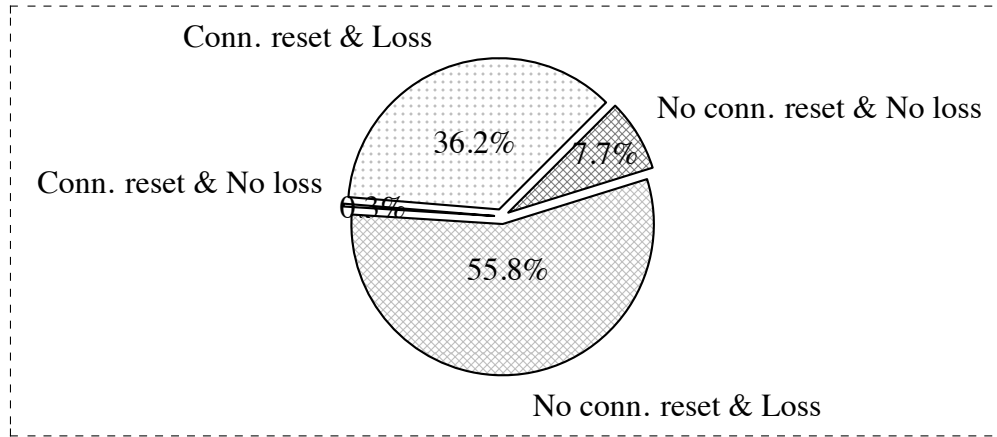


Figure 7.7: The percentages of lossy and non-lossy 5-minute bins for varying RAT with a RAT change split by the presence of connection resets.

no connection resets. The majority of RAT changes, about two thirds, complete without a connection reset. Overall, *inter-RAT handovers are patently lossy and involve a short loss of connectivity in over one third of the bins.*

To further analyse loss related to RAT changes, we identify all distinct RATs present in each 5-minute bin with an inter-RAT handover ¹. Figures 7.8 and 7.9 show the distribution of loss in bins with RAT changes split by the involved RATs. Figure 7.8 shows bins *without* connection resets, while Figure 7.9 shows bins *with* connection resets.

The plots highlight three interesting facts:

1. We only observe *minor difference between Netcom and Telenor in both plots.* This suggest that the same underlying causes lead to this loss in both networks.
2. *The majority of bins with RAT changes include packet loss regardless of the involved RATs.* Loss is, however, much higher when 2G is among the distinct RATs. Bins where 2G is involved occur in poorly covered areas that are characterised by coverage gaps and the dominance of 2G.
3. *Packet loss in bins with connection resets is markedly higher;* loss is over 3% in 90% of the bins. We believe that this loss is a consequence of unsuccessful

¹There can be more than one handover in a 5-minute bin, which means that there is no one to one mapping between the number of distinct RATs and the number of handovers.

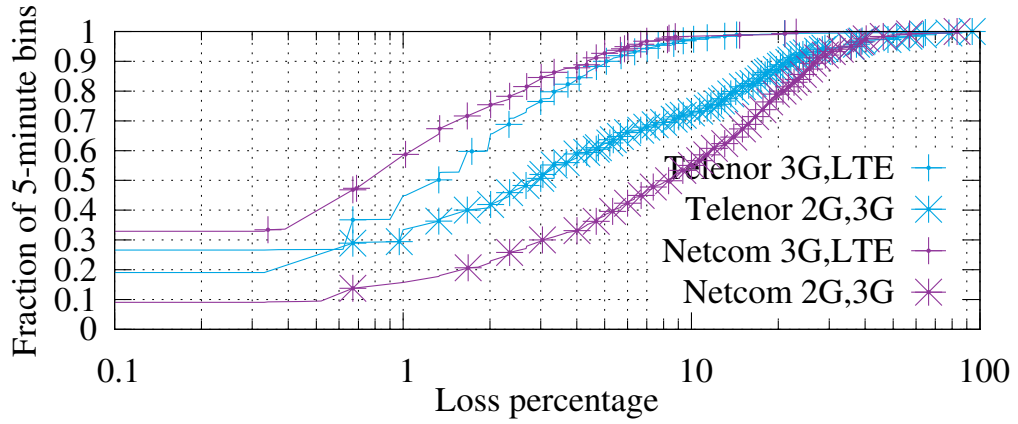


Figure 7.8: Involved RATs for the 5-minute bins with a RAT change and no connection resets. Clear differences between networks. More loss during 2G/3G handovers.

handovers, which initially result in packet loss followed by the connection reset.

To quantify the impact of loss in bins with RAT changes, we look at the loss runs, which we also used in the stationary scenario (see Section 6.2.2). We observe that for the bins without connection resets, most loss runs are of size one. However, loss run size is characteristically different for the bins that involve one or more connection reset. Figure 7.10 shows the probability density function of the loss run size distribution for different RAT combinations for bins with connection resets.

While a sizeable fraction of loss is still random regardless of the network or involved RATs, there are some modes at loss runs of size 5 to 13. Since these modes are present in both networks and in the two sets of RATs, we believe that they are caused by the response of the specific UE to weak or failing coverage.

The loss runs in bins with connection resets consist of two components. First, loss that happens right before the connection reset; we often experience degraded performance before a connection reset. Second, loss that happens between the actual loss of PDP context (EPS bearer) and until PPPd discovers the loss of IP address. We typically detect the loss of connectivity immediately. This detection, however, may take much longer if the modem stops communicating with the PPPd by not responding to PPP echo requests regularly sent by the daemon (PPPd is monitored by usbmodem-listener that is part of NNE and described in Section 4.3.2). These cases can occur when the modem is busy with trying to ex-

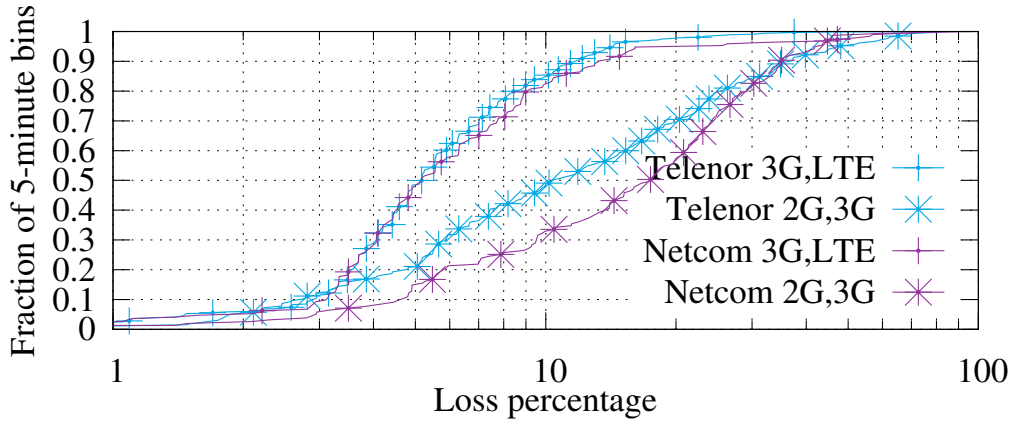


Figure 7.9: Involved RATs for the 5-minute bins with a RAT change and one or more connection reset. Much more loss compared scenario without disconnects, minimal differences between networks, least lossy are bins with 3G/LTE handovers.

change signalling messages with the network. PPPd on the measurement nodes responds to the lack of echo replies by tearing down a stuck connection after six seconds. This partially explains the mode around six in Figure 7.10.

Loss during periods with RAT glitches

In case RAT becomes unavailable, the modems report a special RAT called *No service*. In some cases this RAT is also reported during the inter-RAT handover procedure. *No service* periods mostly happen during the temporary loss of coverage and/or unsuccessful inter-RAT or horizontal handovers. In this subsection, we investigate loss in bins that include strictly one RAT and at least one *No service* period.

As for the periods with inter-RAT handovers, some of the bins with no RAT changes have one or more connection resets. Figure 7.11 shows the split of varying RAT bins with no RAT change according to the presence of loss and connection resets. The overall percentage of lossy bins is slightly smaller (81.6%) compared to the scenario with RAT changes, but it is still high. There are also more (18.8%) non-lossy bins without connection resets compared to the scenario with RAT changes.

Next, we look at loss rate distributions for bins with *No service*. Figures 7.12 and 7.13 show the loss rate for the bins without or with one or more connection

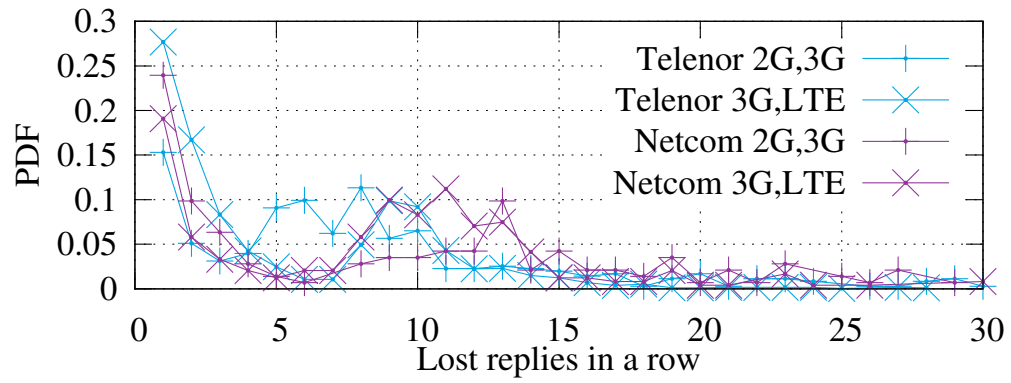


Figure 7.10: Loss runs for involved RATs when there is a RAT change and one or more connection reset. Several modes around 10 packets in both networks. Different modes for different sets of involved RATs.

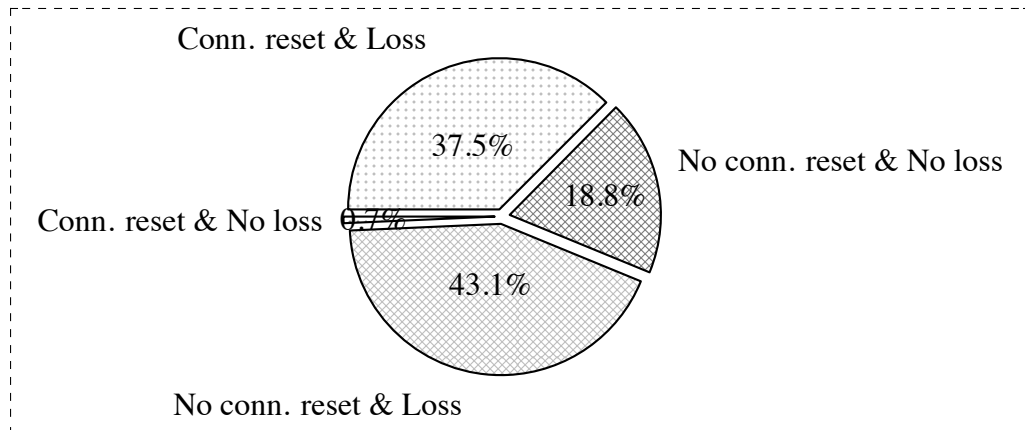


Figure 7.11: The percentages of lossy and non-lossy 5-minute bins for varying RATs with no RAT change split by the presence of connection resets.

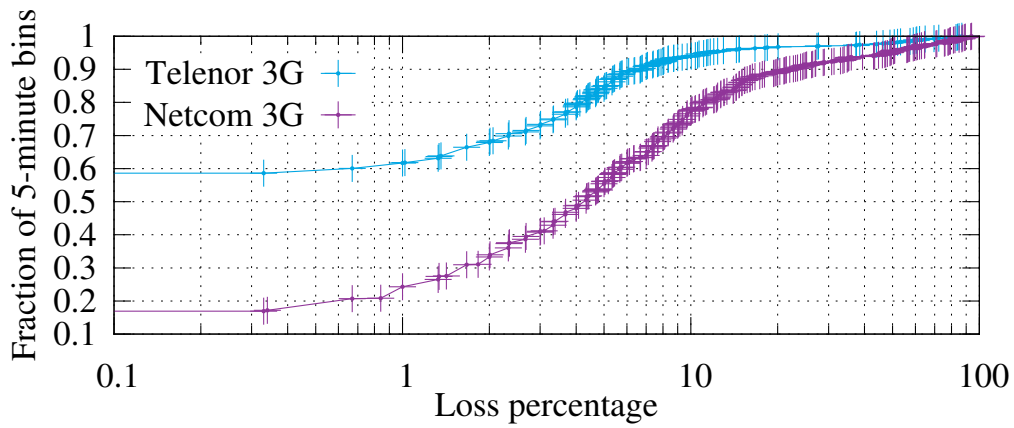


Figure 7.12: Individual varying RATs for the 5-minute bins when RAT does not change and there are no connection resets. Less loss in 3G bins with varying RAT in Telenor compared to Netcom.

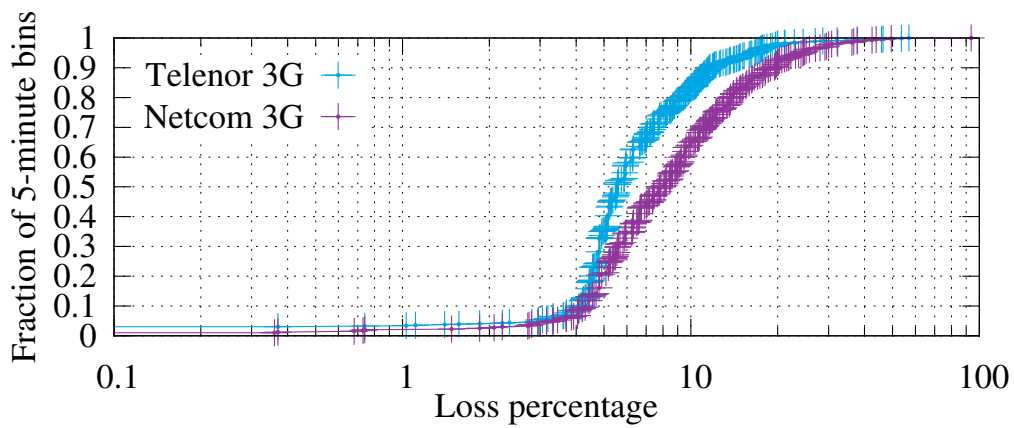


Figure 7.13: Individual varying RATs for the 5-minute bins when RAT does not change and there is one or more connection reset. Differences between networks are minimal, but there is slightly less loss in 3G bins with varying RAT in Telenor compared to Netcom.

resets, respectively.

Here we focus on 3G bins only, since the number of 2G and LTE bins is very low in both networks. This can be explained by the fact that 3G is the dominant RAT country-wide and therefore most handovers happen on 3G. Telenor exhibits much less loss compared to Netcom in bins without connection resets, hinting that coverage problems that lead to *No service* are less prevalent in Telenor. This matches well our out-of-band understanding of the coverage of these two operators; Telenor has a denser deployment of cell towers than Netcom. Loss in bins with connection resets is, however, much higher and very similar across the two networks. We believe this similarity is a product of the non-trivial response of UE to the loss of coverage as explained in Section 7.1.6.

To quantify the impact of loss in bins with RAT changes, we now look at the distribution of loss runs sizes. Figure 7.14 shows the Probability Density Function

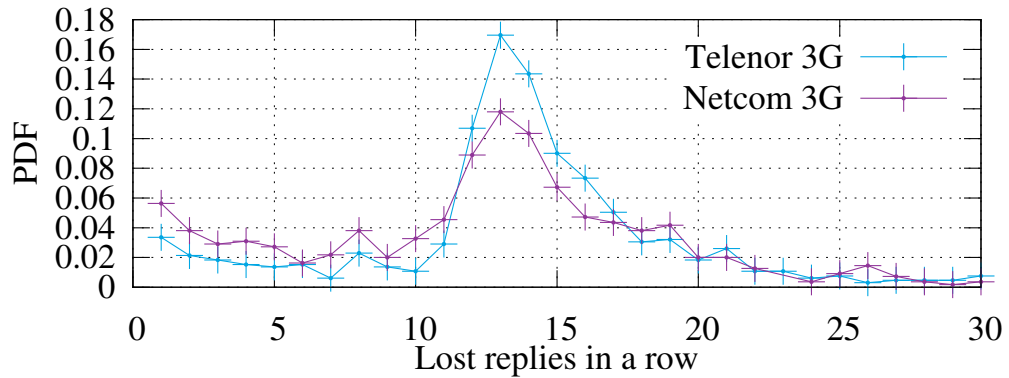


Figure 7.14: Loss runs for individual varying RATs when there is no RAT change, but one or more connection reset in a 5-minute bin intersecting with the loss run. All modes in a shorter range from 11 to 16, with a peak at 13 packets for both networks.

(PDF) of the loss run size distribution for the two networks. These distributions clearly differ from the loss run size distribution for bins with RAT changes and connection resets in two respects. First, there is no random loss. Second, loss run sizes are confined to a narrow range between 10 and 15. These observations indicate that these loss runs must be triggered by temporary lack of coverage followed by the connection resets.

Summary of findings. This section has shown that the loss rates are high in periods with varying RAT, independent of whether there is an actual inter-RAT

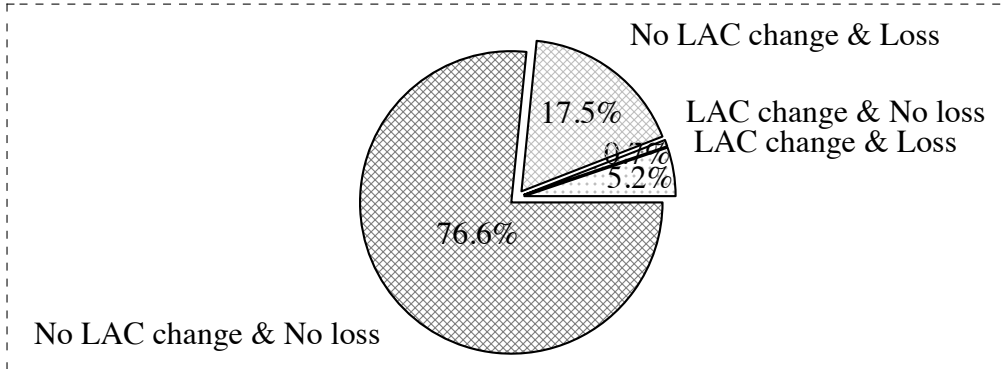


Figure 7.15: The percentages of lossy and non-lossy 5-minute bins for constant RAT split by the LAC change.

handover or not. About 40% of bins with varying RAT also contain a connection reset. If connection resets are involved, we normally also see packet loss, and the loss episodes are more severe.

7.1.7 Constant RAT

This section investigates bins that are characterised by constant RATs (i.e., no inter-RAT handovers), which are located on the leftmost subtree in Figure 7.4. During these periods a connection may experience LAC and cell changes as well as channel quality degradation.

As shown in Figure 7.4, about 30% of packet loss during mobility takes place in bins with *Constant RAT*. Most of this loss, 72%, coincides with changes of serving cells (CID change).

Figure 7.15 divides *Constant RAT* bins based on whether there is a LAC change or not and shows the percentage of bins that fall into four different categories that describe LAC change and loss. The fraction of bins with LAC changes is small (6.3%), which is expected since one LAC mostly covers large geographical areas and LAC changes happens when crossing the boundaries between areas. Connections used in this study experience loss in 88% of the bins with a LAC change. For a smooth handover, the UE needs to be able to communicate with both the current and the candidate cells upon starting the handover procedure. Once the handover is completed, in-flight packets will be re-routed to the new cell. Inter-LAC handovers are slightly more challenging, since they involve additional coordination

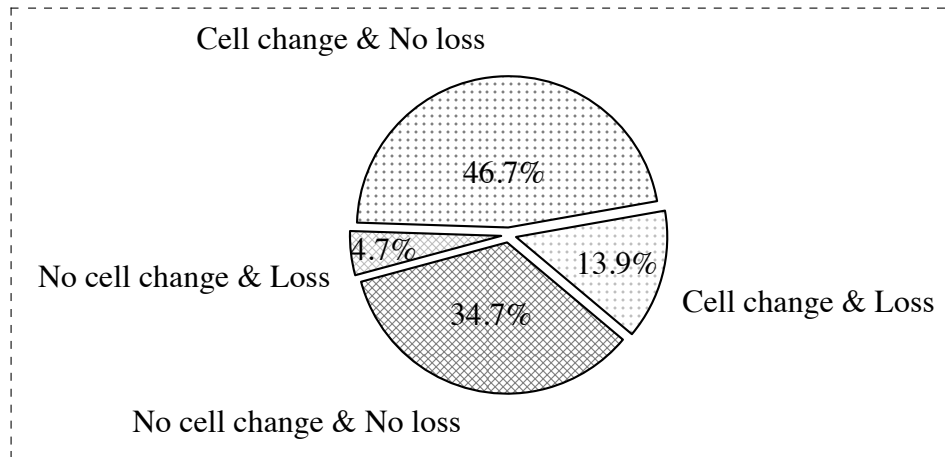


Figure 7.16: The percentages of lossy and non-lossy 5-minute bins for constant RATs split by the LAC and cell changes.

between several RNCs ², i.e., the handover procedure takes longer to complete compared to cell changes within the same LAC.

Section 7.1.7 divides *Constant RAT* bins without LAC changes into four categories that capture both cell changes and loss. The connections experience a cell change in 60% of all bins without LAC changes. These handovers are usually smooth, with 77% completing without a single packet lost. Loss in bins with cell changes is, however, three times higher compared to those without. Figure 7.17 shows loss rate distribution for bins with LAC changes, cell changes, and no changes when the connections are on 3G or LTE.

Loss rates are evidently higher in bins with LAC changes with clear differences between operators. Almost all LAC changes in Telenor involve packet loss, while for Netcom, LAC changes seem to be smooth in 40% of the cases. Hence, this loss appears to be dependent on the network configuration. We also observe that Netcom 3G connections experience significantly higher loss when switching cells compared to Telenor 3G connections. Loss is minimal during LTE cell changes with no clear differences between operators. We believe that loss during handovers can happen due to one of the following three reasons:

1. Short coverage gaps between adjacent cells.

²In theory, an RNC may serve more than one LAC. Private communications with the measured operators confirmed that is not the case in the networks we measure.

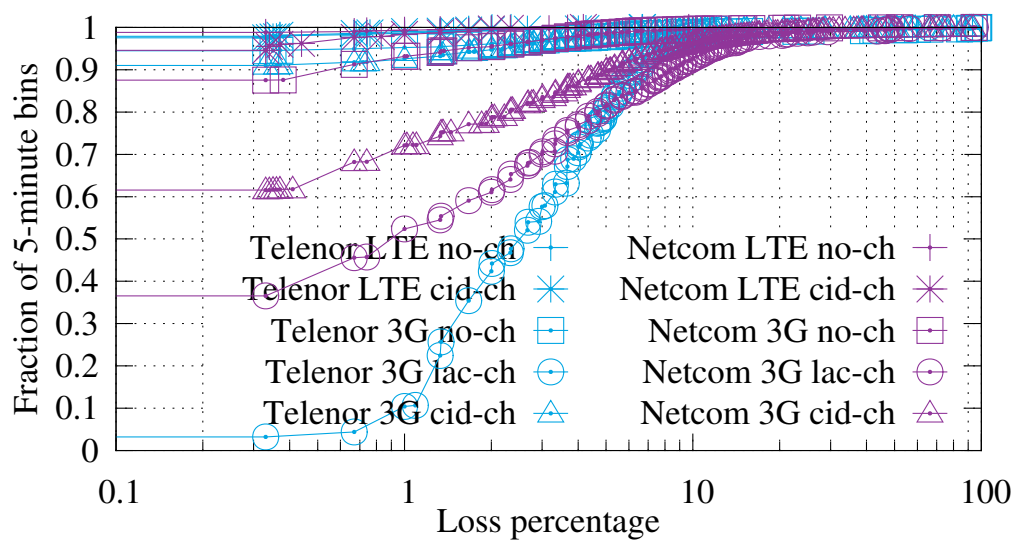


Figure 7.17: Loss rate in a 5-minute bin for 3G and LTE RATs split by the presence of one or more LAC change or cell change, or none of them. In both networks the highest loss occurs when there is a LAC change involved. In Telenor, bins with LAC changes are more lossy compared to Netcom, whereas in Netcom there is much loss during cell changes. When there is no LAC or cell change or the LTE cell changes, the loss is minimal in both networks.

Scenario	$S < 10$	$10 < S < 50$	$50 < S < 100$	$100 < S$
CID-3G	0.48	0.45	0.60	0.72
CID-LTE	0.10	0.21	0.67	NA
LAC-3G	0.85	0.75	0.90	0.79

Table 7.1: Fraction of lossy bins for Netcom cell and LAC changes for different speed (S) categories. The speeds are in km/h.

2. Misconfigured neighbor cell list, which makes affected cells not aware of their neighbors and thus unable to complete handovers successfully.
3. A complex interplay between the timing of the handover decision and trains speed. When deciding to handover, the UE performs an attachment procedure during which it becomes attached to two cells; the current cell and the candidate cell. The handover will break, if the UE loses sight of the old towers during the movement while the procedure is ongoing.

Table 7.1 shows the fraction of 5-minute bins in Netcom that involve loss for different train speed categories and different horizontal handover scenarios. We choose Netcom because it demonstrates significantly more loss during CID changes. The likelihood of experiencing loss during CID changes evidently increases as the speed increases over 50 km/h. LAC changes, however, involve loss independent of the speed, suggesting that the root cause of loss is perhaps related to inter-LAC handover procedure configuration. Note that we have not measured LTE cell changes when the speed is higher than 100 km/h. Trains reach high speeds outside the metro-area and Netcom seem not to have much LTE coverage in these areas.

Summary of findings. Loss is significantly lower in bins where the RAT type is stable. With a stable RAT, cell handovers, and in particular those involving also a LAC handover, is a main cause of loss. There are clear differences in how handovers affect loss between operators.

7.2 Investigating excessive delays under mobility

Packet loss is an important metric to assess reliability of MBB networks. The ability to deliver IP packets end-to-end (no packet loss), however, does not always mean that the packets timely arrive at the destination or back to the source. Further, loss episodes typically precede or supersede a set of packets with higher

RTT values. This is especially prominent under mobility, when inter-RAT and intra-RAT handovers and coverage holes come into play.

In Section 6.4, we investigated delays for the stationary nodes and observed episodes with both large and variable RTTs. In this section, we supplement the delay characteristics with the data from four NNE nodes on long-distance trains that were under mobility most of the time. We measure RTTs under mobility for Telenor and Telia, the two major Norwegian MBB networks. Our goal is to find out whether extreme delay events with the same triangle pattern are also common for mobile connections and identify the possible causes of excessive delays under mobility.

7.2.1 Basic delay statistics

Nodes in motion experienced higher RTTs, likely due to requisite handovers and potential changes in radio access technologies. Figure 7.18 plots the CDFs of average RTTs (left) and maximum RTTs (right) for non-stationary connections on Telenor, classified by RAT and radio resource state.

We observe three differences between static (see Section 6.4.1) and mobile connections. First, mobile connections with mixed RATs no longer have the highest RTT profile. This mixed RAT sample for stationary nodes is anomalous, since handovers are not expected; RAT changes must be due to dynamic cell breathing, over-crowded or rural cells, and other network infrastructure problems that are not the focus of our study. Second, 3G-DCH connections exhibit significantly (not just slightly) higher delay than LTE. Third, a large fraction (40%) of 3G-DCH and 3G-FACH bins suffer from maximum delays higher than one second, with approximately one third of LTE bins showing maximum RTTs over 100msec. Perhaps most interesting, however, is the high variability in RTTs we observe for static and moving 3G-DCH and LTE connections, which we investigate next.

7.2.2 Frequency of triangle events

Table 7.2 shows the maximum and minimum hourly rates of triangle events for our four mobile nodes for LTE, 3G-FACH, and 3G-DCH. With the exception of LTE, mobile node results are comparable to the worst 5% of stationary connections. Although our mobile sample is small, results for both operators and all RAT/RRC combinations (including those not shown in the table) are consistent: triangle events are more likely when nodes are moving.

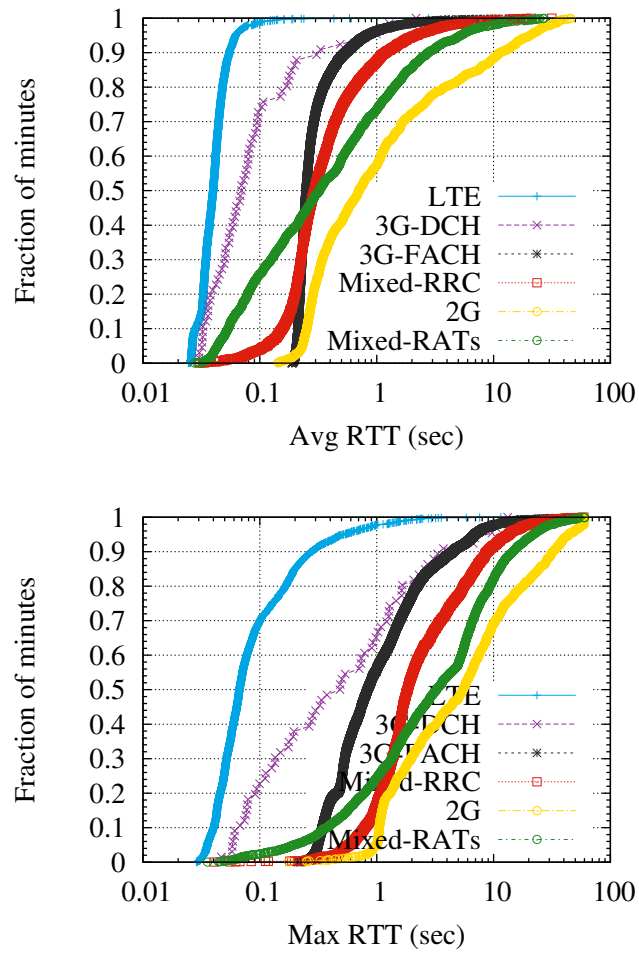


Figure 7.18: Average and max RTT for mobile connections on Telenor (one-minute bins). 40% of 3G-DCH and 3G-FACH bins exhibit maximum delays ≥ 1 second.

Table 7.2: Triangle event frequency for mobile nodes.

Category	LTE	3G-DCH	3G-FACH
Telenor			
max hourly rate	1.18	32.2	13.01
min hourly rate	0.4	4.16	5.65
Netcom			
max hourly rate	0.72	8.51	18.51
min hourly rate	0.25	2.3	7.39

7.2.3 Summary of findings

This section is a supplement to a large-scale measurement study of round-trip delays in stationary nodes (Section 6.4) with the measurement data from four mobile NNE nodes. The observed extreme, multi-second delay episodes were more common when our mobile nodes were actually in motion. We also observe that maximum RTTs can reach several seconds, especially when nodes are moving. We observed that triangle duration distributions for mobile nodes are comparable to stationary nodes.

In summary, triangle events are fairly common when connections are moving, and increase in frequency when the signal is poor. Their duration varies depending on the underlying change. We also spot differences between operators that reveal room for network improvement via tweaks to state transition and handover timers. Further, we observe triangle events on both the uplink and downlink, albeit more frequently on the former.

7.3 Summary

This chapter has analysed the causes of loss and round trip delays in MBB networks under mobility. In general, we established that loss in MBB is significant in a mobile scenario, and much higher than in the stationary case, presented in Chapter 6. We further observed that extreme delay episodes lasting several seconds to be more common when nodes are in motion.

Packet loss. Disturbances or handover between different RATs is a main cause of loss, accounting for about 70% of the total. Such RAT changes also often involve a reset of the data connection between the UE and the network, which

mostly involves heavy packet loss. Cell changes are also an important source of loss, and cell changes that also involve a LAC change are the worst.

The observed dominance of loss during RAT changes highlights such handovers as an area that warrant particular attention from mobile operators. The inter-RAT handover procedure is complex, and involves interaction between the UE, the RAN and the CN. The most efficient way to reduce packet loss is to improve the procedures for how such handovers are performed. The number of such handovers should be limited, and packets in transit should be buffered or retransmitted to avoid loss.

There are significant differences between the two networks measured in this study with respect to loss during cell changes. While Telenor experiences significantly more loss during LAC changes, Netcom sees more loss during normal cell changes. These differences indicate that operators still have a significant potential for reducing loss through better configuration settings in their network.

To verify some of our findings, we conducted a drive test in Oslo area by placing the measurement node in a car. In total, we have collected over 5 hours of measurements for the two networks. These measurements confirmed that in Telenor, almost all (92%) packets were lost during the periods with varying RAT. In Netcom, around one half of loss happened in bins with varying RAT too, while the second half was during periods with cell or LAC changes. As expected, we have not observed any case with varying RAT and a temporary loss of service, as it is very unlikely to have coverage holes in the city. In other words, results from the drive test confirm and highlight that inter-RAT handovers are prone to high packet loss even in well covered areas.

Excessive delays. We observed that triangle duration distributions for mobile nodes are comparable to stationary nodes. We also found that excessive delay episodes, which we call triangle events due to their triangular-like pattern, increase in frequency when the signal is poor. In addition to the RRC state transitions affecting the delays of both stationary and mobile nodes, the latter experienced more triangle events when handovers were ongoing. We have however observed differences in delays during handovers between the two operators we measured. We also identified that most of buffering during handovers occurs on the uplink.

End-to-end measurements used in this study are useful for quantifying and characterising the problem. However, to localise the root causes of packet loss or

extreme delays, this might not always be sufficient. Network side data or measurements from the RAN could give more insights into the potential causes and assist in improving the network.

Chapter 8

Evolution of MBB networks

One of the mandates related to the NNE testbed is an annual report on MBB reliability in Norway. As of February 2017, there were three such reports published, for year 2013, 2014, and 2015, and the report for year 2016 was being prepared. The measurement data used to produce the reports is the same used in Part II of this thesis. The reports also contain a set of additional measurements and longitudinal comparison of the different aspects of MBB reliability and performance.

Different studies presented in Part II cover different measurement time periods ranging from one month to more than one year, as described in Section 5.3. In order to be able to capture the evolution of reliability and performance over time, it is important to use the same set of metrics for the entire measurement period. In this chapter, we look at how the reliability and performance of Norwegian MBB networks have evolved over time from July 2013 until December 2015. We also explain some of the causes that led to the degraded packet data service of individual networks.

8.1 Shorter and less frequent downtimes

Figure 8.1 shows how the fraction of connections with MTBF less than 1 day has developed throughout our measurement period. The period is divided into five periods of six months each. For the majority of periods and operators, the fraction of connections with $\text{MTBF} < 1$ day was stable, around 20%. The reason for the higher fraction of connections for Telenor in the second half of 2013 and the first half of 2014 is related to the RRC state machine configuration problems (discussed in Section 6.3.7), which was resolved in June 2014. Network Norway

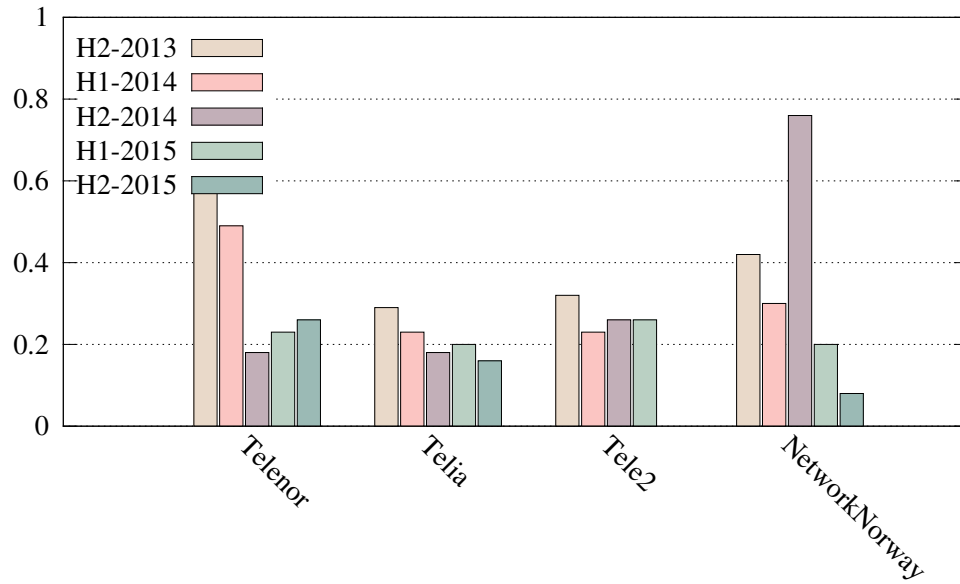


Figure 8.1: The fraction of connections with MTBF < 1 day.

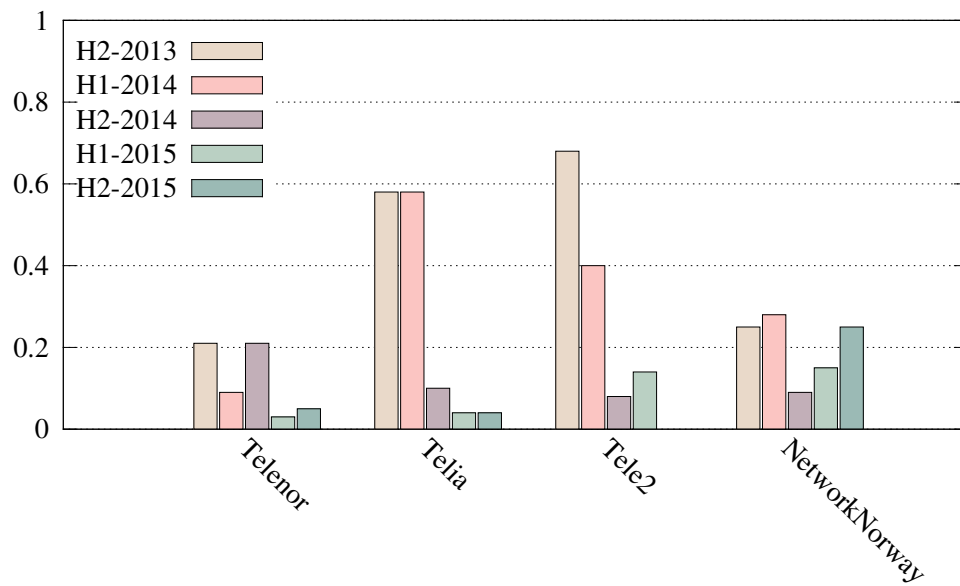


Figure 8.2: The fraction of connections with MTTR > 5 minutes.

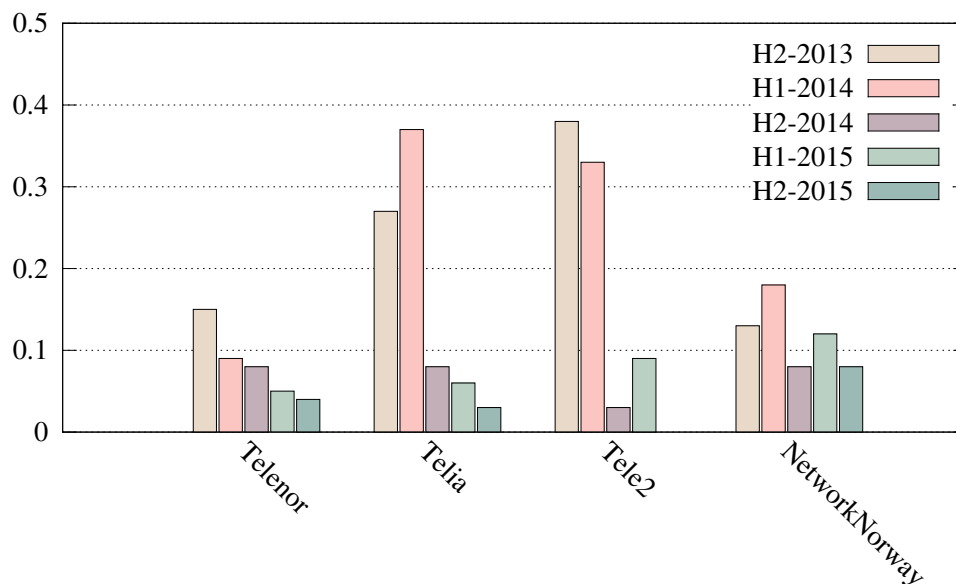


Figure 8.3: The fraction of connections with the average downtime > 10 minute per day.

had major issues with its CN in the beginning of November 2014, when almost all connections were terminated multiple times. Network Norway later confirmed that this was caused by the configuration problems in one of their CN elements. This explains the increased MTBF fraction in 2014, which was no longer the case for 2015.

Figure 8.2 shows how the fraction of connections with MTTR of more than 5 minutes developed throughout the measurement period. Similar to the MTBF, the fraction of long-lasting connection failures has been reduced over the years for most of the operators. Network Norway experienced a set of major outages in 2015, when many connections were terminated and could not recover for three hours. First such outage was observed in March 2015, after customer migration from Mobile Norway to Telia's (Netcom's) RAN. The problem, sporadically affecting different connections at different times, persistent until the end of 2015. We do not know the root cause of this problem, but the downtime duration of exactly three hours suggests that some timer must expire before the network can accept the re-establishment of the connection. For a typical user such problem could only be resolved by restarting the mobile device. For an M2M device such a failure can result in a complete halt until the device is power-cycled. The described problem

is the main reason behind the increased MTTR for Network Norway, especially in the second half of 2015.

Figure 8.3 shows how the fraction of connections with the average downtime of more than 10 minutes per day has developed throughout our measurement period. The downtime is indirectly proportional to both MTTR and MTBF: the shorter it takes to re-establish the connection (MTTR) and the longer is the time between the two consequent connection failures (MTBF), the less likely to have an accumulated downtime of 10 minutes or more per day. The figure therefore follows the same pattern of an overall improvement as in Figures 8.1 and 8.2. What also observe that the downtime distribution inherits the pattern of the MTTR and MTBF distributions. This can be explained by the fact that even a single, but long-lasting (longer than 5 minutes) connection drop per day directly contributes to the average downtime, shown in Figure 8.3.

In conclusion, we have observed reduction in downtime from 2013 until 2014, and especially in 2015. Less than 10% of Telenor, Telia (Netcom) and Tele2 connections had more than 10 minutes downtime per day in 2015. In the second half of 2015, the fraction of Telenor and Telia (Netcom) connections with 10 minutes downtime per day was 4% and 3%, respectively. None of the measured Tele2 connections had more than 10 minutes downtime per day in the second half of 2015¹.

8.2 Significant reduction in loss

We have observed reduction in packet loss rates across all operators from 2013 until 2014. In 2015, loss rates continued to decrease for Telenor and Telia (Netcom) connections, which was not always the case for Tele2 and Network Norway connections. Figure 8.4 shows how the loss rates developed throughout our measurement period for Telenor and Telia (Netcom). The figure is similar to Figure 6.18 and plots the daily median loss rate per day for all connections. We observe that with a few exceptional days and the problem in Telenor's RAN until July 2014 discussed in Section 6.3.7), packet loss was low for both networks, especially in the second half of 2015.

Figure 8.5 shows the median packet loss for the four networks throughout the measurement period, divided into the periods of six months. The figure shows

¹ The number of Tele2 connections in the NNE testbed was greatly reduced in the second half of 2015 as shown in Figure 4.7. We have therefore not included results for Tele2 in this period.

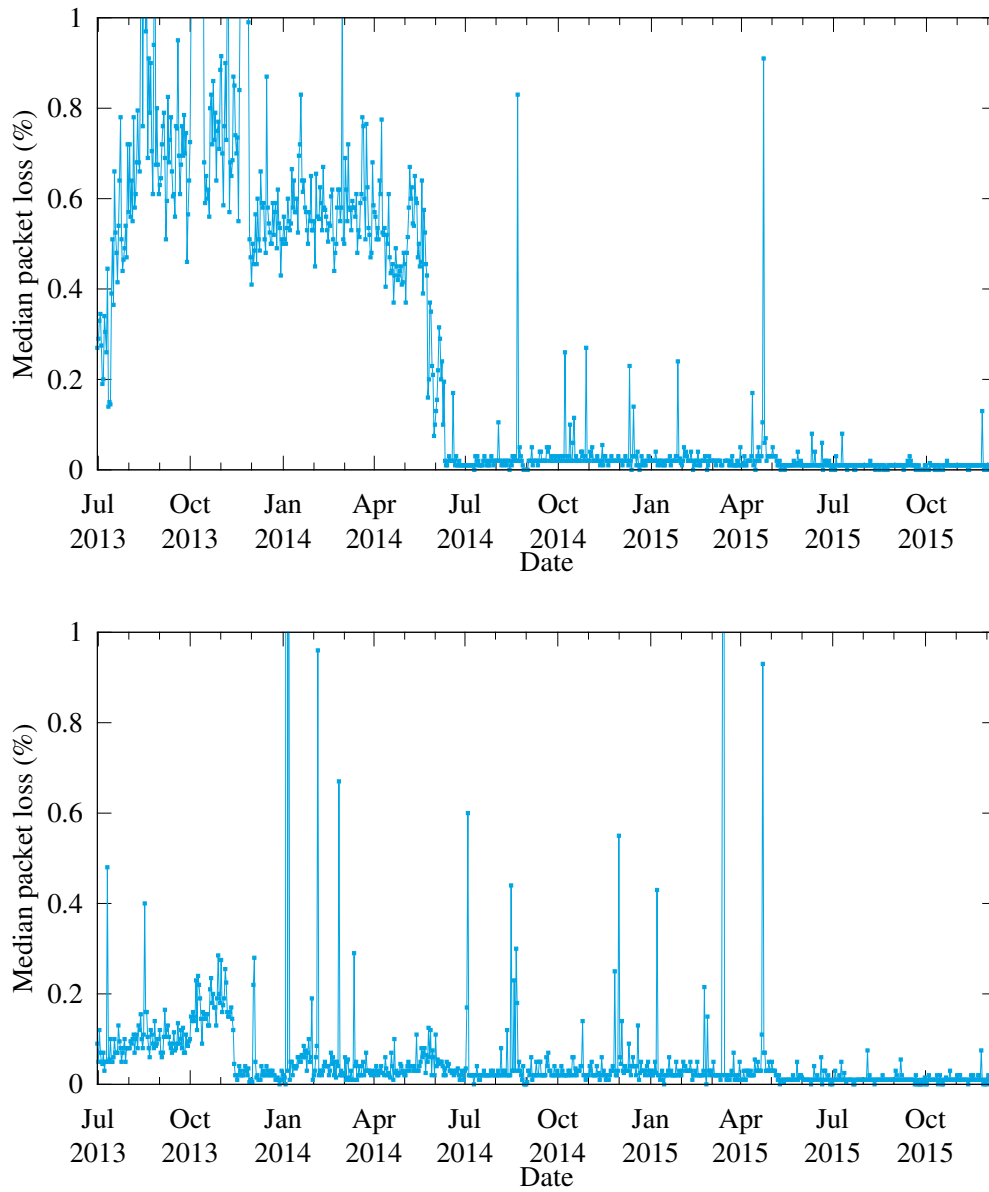


Figure 8.4: Median loss rate evolution in Telenor (top) and Telia (Netcom) (bottom) over time.

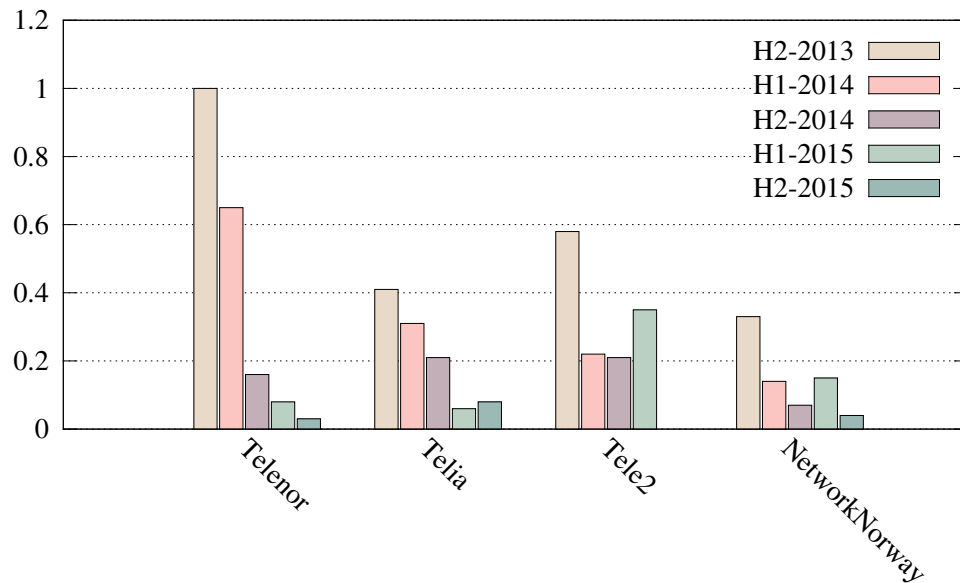


Figure 8.5: Median loss rate evolution.

the reduction in packet loss for both Telenor and Telia (Netcom) in 2015 when compared to the previous years. For Tele2 and Network Norway there are slight changes in loss rates from 2014 until 2015, but Network Norway clearly had lower loss rates in the second half of 2015. For both of these networks we observed a reduction in packet loss after their connections were transferred to Telia's (Netcom's) RAN in March 2015. It is also important to mention the development of LTE in 2014 and 2015 all over Norway, which contributed to the decreased loss rates compared to the 3G.

8.3 Fewer and less severe large failure events

Figure 8.6 shows large failure events during the entire measurement period. The detection of the events shown are based on the same principles as for the Figure 6.9 in Section 6.2.3, except that for each connection, in addition to packet loss, connection downtime is also taken into account. As seen from the figure, there are fewer large failure events in 2014 when compared to 2013. Further, in 2015 there were no large failure events with severe loss rates (severity is proportional to the diameter of the circle). There were, however a set of events in 2015 affecting 100% of our connections from a particular MBB network. Another observation

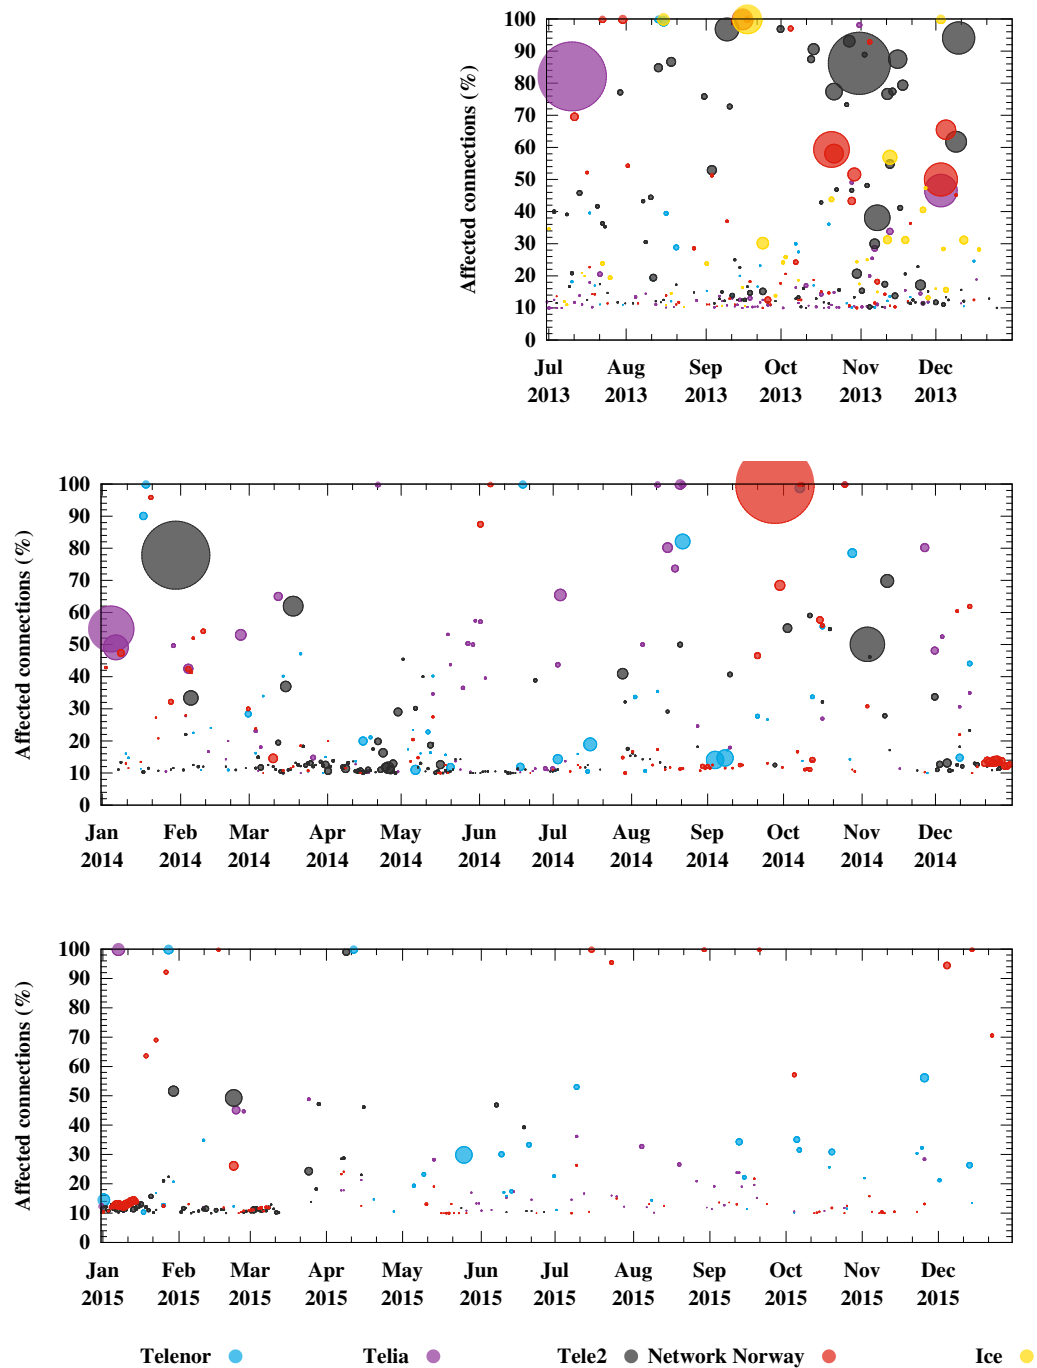


Figure 8.6: Large failure events from July 2013 until December 2015.

is that the number of events in Tele2 and Network Norway networks has reduced after their migration to Telia's (Netcom's) RAN in March 2015.

8.4 Summary

This chapter has shown that the reliability of MBB networks has improved significantly throughout our measurement from July 2013 until December 2015. The overall stability of the measured MBB networks has improved over time. In 2015, when compared to 2013 and 2014, connections had shorter downtimes, lower overall packet loss rates and fewer failure events affecting individual operators.

We see the three major contributors to the improved performance of Norwegian MBB networks over the last couple of years. First, an ongoing LTE deployment from 2014 has contributed to lower loss rates and shorter connection downtimes thanks to a simplified state machine, hybrid ARQ, the use of turbo codes for high-performance forward error correction, and other improvements. Some of these improvements were also backported to HSPA and HSPA+ networks. Second, we observe improvements in coverage. In particular, in 2015 there were fewer stationary NNE nodes with poor signal quality compared to the previous years. The same is for the coverage holes experienced by our mobile nodes, which became less in the recent years. Finally, we believe that our long-term measurements and studies also played an important role for revealing certain performance problems. Every annual report we published has received attention from the public and from the operators, in some cases leading to concrete actions, such as reconfiguration of RAN at scale. The majority of measurements collected from NNE nodes as well as annual reports on Norwegian MBB availability are publicly available.

Part III

Towards more reliable and flexible networks

Chapter 9

Routing limitations in current LTE networks

9.1 The problem of inflexible routing

In the previous parts we looked at different aspects of reliability in MBB networks. One of our findings was the fact that some packet loss happens beyond the RAN, in between the base station and the CN, in the CN itself or in between the CN and the end host. We also found that reliability can be increased through multi-homing, assuming that a user has connectivity to more than one MBB network. Given the built-in geographical diversity of MBB networks, there is a potential for multi-homing and hence increased reliability while being connected to a single MBB network. Current LTE standards allow a UE to establish multiple simultaneous bearers towards distinct PDNs. These PDNs can be served by packet gateways with physical locations different from the centralised CN. This can enable users to access the same end-host over diverse physical paths and hence increase the reliability.

However, the majority of the LTE networks today are still solely imposing and executing routing decisions and almost all connectivity management, charging, and traffic policing in the mobile packet core. IP payloads are tunnelled all the way from base stations to the CN before any forwarding or routing decision can take place. Traffic tunnelling prevents mobile networks from leveraging the flexibility of the hop-by-hop routing enjoyed by the rest of the Internet, thus resulting in suboptimal routing that makes it difficult to offload the mobile packet core. Breaking this rigid routing will spare traffic exchanged between devices within

the same mobile network from a detour via the mobile packet core.

A more flexible routing scheme is likely to become critical in next generation 5G networks. Mobile data traffic is projected to grow almost eightfold between 2015 and 2020 [24]. Further, the rise of machine-to-machine communication and IoT will result in traffic patterns that are local in geographical scope [20]. These changes fundamentally challenge some of the well-rooted assumptions about user mobility and traffic patterns. For facing these challenges, we believe that routing in mobile networks would benefit greatly from assimilating to pure IP networks which have an impressive track record in efficiently adapting to an ever-changing use cases and traffic patterns.

In this part, we propose MULTEX – an opportunistic step towards a packet switched model for current and future generation mobile networks. We intend by MULTEX to show that existing standard features in today’s LTE networks can be leveraged to break the current monolithic architecture of mobile networks. The major goal of MULTEX is to introduce diversity to the way mobile users access packet data service. MULTEX transforms a UE into a multi-homed device that is simultaneously and independently connected to several packet data networks (PDNs) with different geographical or logical boundaries. First, MULTEX gives the UE more control over IP routing decisions. This allows users to alternate the path for their IP flows among the available PDNs, enabling simultaneous access to different private and public networks. This gives increased reliability and QoE. Second, depending on how much content and communication can be moved closer to the edge, earlier termination points will markedly offload traffic from EPC. Finally, MULTEX relies on support for multiple simultaneous PDN connections, which is a standard LTE feature, mandated by recent 3GPP releases and supported by the majority of LTE networks and devices in the market. MULTEX can therefore be realised on existing infrastructure with no modifications to the existing signalling protocols. By relying on a standard LTE feature, MULTEX ensures the added flexible routing and forwarding does not come at a price of adding extra transparent middleboxes like in similar proposals [22, 99].

9.2 The support for multiple PDN connectivity in current LTE standards

9.2.1 EPS bearers and tunnels

A UE that is attached to an LTE network has a default EPS bearer established at all times. The default EPS bearer can have one or more dedicated EPS bearers associated with it and each EPS bearer is assigned with different QCI. The dedicated EPS bearer can be assigned a guaranteed bit rate, whereas the default EPS bearer is always best effort. While a UE can be simultaneously attached to up to 8 EPS bearers [6], the exposure of respective network interfaces to the users is not common. Once an EPS bearer is established, two tunnels, namely S5 and S11, remain established throughout its lifecycle. Depending on the RRC state, two additional bearers and respective signalling connections, namely Data Radio Bearer (DRB) and S1, have to be re-established or released. Concatenation of DRB and S1 bearers is referred to as E-UTRAN Radio Access Bearer (E-RAB).

9.2.2 Traffic offloading

To address the demand of off-loading certain traffic from the EPC, 3GPP Release 10 has introduced solutions for Local IP Access (LIPA) and Selected IP Traffic Offloading (SIPTO) [5]. Both solutions are primarily targeted for Home eNBs (HeNBs) – small cells, typically found in indoor premises and connected to both mobile radio and corporate or local network. LIPAs allows users to access their local networks through a local APN, while SIPTOs enables to offload selected IP traffic from the EPC. Both LIPAs and SIPTOs require a local Local PGW (L-PGW) to be collocated or deployed close to the (H)eNB and rely on additional eNB functionality to split the traffic to local and non-local and then route it accordingly. The two solutions also introduce certain modifications to the signalling protocols.

9.2.3 Multiple PDNs

A UE can be simultaneously attached to up to 11 EPS bearers [6]. However, there are only eight possible logical data channel identifiers in the RRC protocol, thus limiting the possible number of simultaneous default or dedicated EPS bearers to eight. The most common use of multiple EPS bearers is the VoLTE service that mandates a separate IMS PDN with one default and one dedicated EPS bearer to

be used for signalling and streaming, respectively. Except for VoLTE, while being standardised and technically supported by most LTE networks and UEs, exposing multiple default EPS bearers to the end users is not common.

9.3 Shortcomings with the mobile network architecture

Historically, unreliable transmission between RAN and CN has motivated the tunnelling of data traffic between base stations and respective core gateways. Transport medium reliability has improved significantly over the past decade with extensive deployment of fibre to interconnect the RAN and the core. Tunnelling, however, has become the status quo and has had a significant impact on the architectural evolution of mobile networks. This is the basis for numerous constraints that remain present in LTE networks and create barriers for certain features and functions commonly found in traditional IP networks. Next, we present several shortcomings with the current mobile network architecture that are direct consequences of tunnelling.

9.3.1 Suboptimal routing due to centralisation

A centralised architecture makes it easier to implement traffic engineering, access control and accounting. However, the main drawback of centralisation is the fact that all IP packets must first reach the EPC before they are forwarded to their final destinations. This stretches end-to-end delay and imposes a significant load on the EPC. The situation is the opposite in traditional IP networks, where IP packets traverse through one or more intermediate routers to reach the destination. Routers always try to find the most optimal route and hence reduce the number of hops to the minimum. This is not the case in LTE networks, where even the IP flows between the two UEs attached to the same eNB need to go through the EPC.

9.3.2 Statefulness

Once an EPS bearer is established, the state of the GTP tunnel is maintained across multiple EPS nodes. Additionally, separate tunnels and states exist for the uplink and downlink channels. State management helps to save radio and other network

resources, but creates extra complexity and increases end-to-end delays due to different timers of the state machines and transition between states.

9.3.3 Lack of geographically limited networks

Some LTE operators offer private APNs to their corporate customers, thus realising a VPN concept by the means of a separate PDN. However, private PDNs suffer from the same suboptimal routing problem due to tunnelling. Their IP packets are decapsulated at the EPC first and then encapsulated again into a VPN tunnel towards a company. SIPTO solution provides one way to solve the problem by deploying a specialised (H)eNB at a company, limiting the scalability of such approach. Further, since mobile networks are unique in their geographical awareness, the concept of virtual private or public networks could be extended for private customers. For example, users residing in the same geographical area be offered with a separate PDN that is unique in features, not offered by the standard PDN for Internet connectivity. One of such features could be native IP broadcast and multicast services. This would open doors to novel services and applications, such as peer-to-peer discovery, emergency alerts, and targeted ads.

9.3.4 Lack of reliability through built-in diversity

Even in centralised LTE networks, there are typically several S/P-GWs that are serving users from different geographical areas and/or connecting them to different PDNs. Assuming that a user can be connected to two different PDNs over two independent SGWs and PGWs, only a part of the entire physical path would be shared for the individual IP flows. A sizeable number of applications that demand path diversity could benefit from increased reliability through multi-homing. For example, a mobile payment terminal could be simultaneously connected to the two distinct PDNs and implement failover mechanisms on top of the corresponding network interfaces.

9.4 Related work

Over the past few years, there has been a number of initiatives that aim to address some of the limitations imposed by the centralised mobile network architecture. In this section, we provide an overview of the most relevant initiatives.

9.4.1 LIPA and SIPTO

Several works studied the applicability and gains of LIPA and SIPTO. Requirements to support LIPA/SIPTO were discussed by Samdanis et al. [82] along with a DNS-based approach for fine-grained offload control of LIPA/SIPTO traffic. The authors also investigate data offloading issues with LIPA/SIPTO with respect to network deployment focusing on capacity, energy efficiency and improvement of radio propagation. Eido et al. [28] used SIPTO to selectively offload mobile IP traffic in order to use servers deployed within the metro network at or above the RAN. They show that in the next five years as much as 30% of backbone traffic can be offloaded by distributing servers for mobile traffic.

In both LIPA and SIPTO, traffic splitting is done by either the means of packet inspection at the (H)eNB or by requesting UEs to use separate PDN connections. The former demands powerful (H)eNB hardware while the latter, like MULTEX, requires UEs to support multiple PDN connections. However, the use of multiple default EPS bearers, the key component of MULTEX, is not thoroughly discussed and in particular lacks details about routing decisions. Additionally, changes are needed in the network for the L-PGW and L-SGW resolution as well as to inform UEs about the support and presence of LIPA or SIPTO.

9.4.2 NFV and SDN based initiatives

Recently ETSI formed NFV industry specification group [69], which has already attracted attention of mobile network vendors and operators. NFV, currently being adopted in both experimental and commercial environments, proposes the mobile network components and functions to be realised as software instances on commodity servers or in data centres. The approach heavily builds on Software-Defined Networking (SDN) principles that include virtualised networks components with service orchestration and a centralised control plane. SDN and NFV promote flatter networks, where the placement of mobility anchor points is closer to the UEs. This way traffic can be offloaded from the core network by the means of local breakout.

An SDN based EPC architecture that centralises the control plane functionality and reduces signalling was proposed by [23]. While promising, their solution requires changes to eNBs to be able to connect to Open Flow controllers that reside in EPC. SoftRAN, a software defined centralised control plane for LTE RAN [38], abstracts all base stations in a geographical area as a virtual station that consists of a central controller and physical radio elements. SoftRAN develops the concept

of a local geographical network managed by a centralised controller. MULTEX does not rely on the presence of the controller and instead promotes geographically limited networks throughout multiple PDNs. In [75], the authors compare the architectural principles of EPS with SDN and propose an evolutionary path for the former, which is based on traffic flow templates, decouples control and user planes and addresses multi-homing, off-loading and route optimisation issues. Although the idea behind the proposed scheme is similar to MULTEX, their solution, however, requires introducing additional states and Traffic Flow Templates (TFTs) at each data plane node. SDN-based distributed mobility management with distributed PGWs and a centralised control plane [54] suggest a route optimisation strategy for internal traffic exchanged between LTE users. The work relies on the IP address prefixes to separate device-to-device traffic, whereas MULTEX does it transparently. A model for the placement and deployment strategies of different network functions in an NFV/SDN environment is presented in [18]. MULTEX does not impose restrictions on how virtualised gateways should be placed in the mobile network, but would benefit from reduced processing delays in the decomposed gateway implementation proposed by the authors.

9.4.3 Other initiatives

Several notable works proposed different distributed architecture models for LTE control and planes [49, 53, 65]. MULTEX builds on the distributed architecture with L-PGWs, but at the same time leverages from the fact that an LTE UE natively supports multiple PDN connections and could make flexible routing decisions.

Chapter 10

MULTEX: towards more flexible routing and services in mobile networks

In this chapter, we present MULTEX – a model for more flexible routing and services in current and future mobile networks. We first discuss the motivation for MULTEX along with design goals and required changes. We continue with potential use cases that could benefit from MULTEX. Section 10.5 discusses the road towards implementing MULTEX and describes possible mechanisms for PDN discovery, routing and decision models. Section 10.6 shows a possible implementation of MULTEX in an existing LTE network. Section 10.7 evaluates key aspects of MULTEX. We conclude in Section 10.8.

10.1 Motivation

LTE is an all-IP network, but due to tunnelling remaining to be the central part of it, it is difficult to leverage basic routing mechanisms that are commonly found in IP networks. Introducing multiple local gateways makes it possible to divide the previously large and monolithic network into multiple smaller networks that are grouped based on physical, logical or geographical boundaries and principles. Exposing those gateways to be used by UEs opens up for many new services and applications.

Even though network function virtualisation initiatives are gaining wider adoption among operators that are gradually decentralising their networks, the concept

and potential of multiple simultaneous PDN connections is not well developed. While the chipsets of most LTE UEs support simultaneous PDN connections, control over their management is normally not exposed to the user. More specifically, while some LTE modems already export more than one network interface that can be used simultaneously by the OS, this is almost never the case for smartphones, or their operating systems, such as iOS and Android. Multiple PDN connections is a standard feature of LTE allowing up to eight simultaneous EPS bearers, as described in Section 9.2. We believe that the limited exposure of the feature that is already supported by any VoLTE-capable UE, is a missed opportunity and deserves to be unveiled.

The use of multiple PDN connections would increase flexibility in routing decisions and shift the decision-making process to the UE. Being simultaneously connected to several PDNs means that a UE has access to multiple network interfaces with unique IP addresses. This in turn allows a UE to make informed routing decisions based on different criteria. For the uplink, the flexibility in routing would mean that a single IP destination might be reached over more than one PDN connection. For the downlink, depending on the network security policies, a single UE becomes reachable over multiple PDN connections.

10.2 Design and philosophy

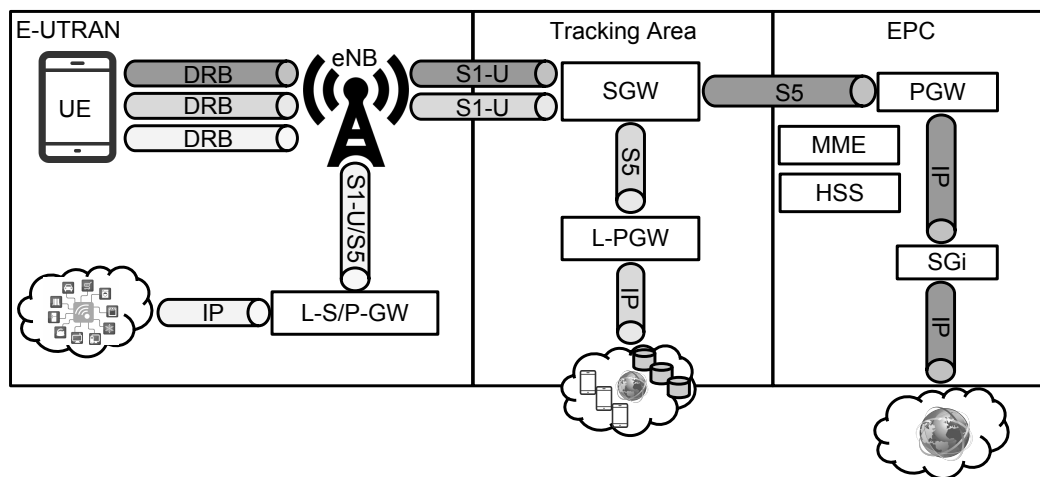


Figure 10.1: Conceptual architecture of MULTEX.

A key design goal of MULTEX is simplicity and pragmatism, avoiding the

extra complexity. MULTEX does not require any major change in the network in order to function, but relies on mechanisms and functions that are well defined and operational in today's LTE networks. The main idea behind MULTEX is to bring more diversity in how packet data service is accessed and used, while maintaining the high degree of backwards compatibility.

Figure 10.1 shows the conceptual model of MULTEX, illustrated by a conventional scenario when a UE is simultaneously attached to three different PDNs that are terminated at different network layers, namely eNB, TA and EPC.

MULTEX builds on PDN connections are already part of LTE standards and supported by the majority of UEs and operational LTE networks. The trend of deploying S/L-GWs closer to the edge is also becoming prevalent. MULTEX therefore requires only a minimal number changes in the network. We believe that MULTEX can be as a stepping stone towards a decentralised, connectionless and pure all-IP mobile network.

10.3 Required changes

The major element MULTEX requires is a mechanism to communicate available PDNs to the UEs, and some additional logic at UEs for EPS bearer establishment and routing. We discuss different possibilities for implementing these mechanisms in Section 10.5. In the following, we discuss the functional blocks of different network components that MULTEX relies on.

UE. Most LTE chipsets support multiple PDN connections, but their management is rarely exposed to the operating system via Radio Interface Layer (RIL) or other means. Implicit activation is done automatically if IMS or other services are used and there are separate APNs configured with the corresponding APN types. MULTEX demands a way to discover the available PDNs, to establish and use additional default EPS bearers, as well as the ability to directly or indirectly manipulate the routing table.

RAN. Strategies and mechanisms for deploying L-S/P-GWs collocated or separately at eNB, within the TA or elsewhere, have been proposed by different works, discussed in Section 9.4. MULTEX does not impose any constraints on how and where in the network breakpoints within the EPS should be realised. We believe that operators should adjust their strategies based on the actual traffic patterns and the number of users in different regions. It is up to the operator to define the best strategy for decentralising and diversifying the network in a way that adds the most

value to its users.

EPC. We expect that in most MULTEX usage scenarios the announcement of available PDNs will be initiated by the EPC. Another necessary element is the SGW resolution procedure done by MME. By default MME resolves SGW for a particular UE based on Tracking Area Identity (TAI) that it receives from the eNB. In some cases this might not be feasible, for example, if a SGW is deployed at the eNB. Since a network operator fully controls the MME as well as the DNS servers, it should not be difficult to implement the necessary SGW resolution schemes based on the APN, TAI, eNB identifier or other contextual attributes.

10.4 Use cases

MULTEX can bring value in a great variety of use cases. We enumerate only a few of them that we believe have the largest potential.

Content delivery and distribution. Provided the network has content cache servers collocated with or deployed close to L-PGWs, separate PDNs could be created per (group of) content or service providers. This might be applicable both to traditional Content Delivery Networks (CDNs) as well as specialised content types, such as firmware updates over the air.

Virtual private networks. It is becoming increasingly popular to let an employer take control of (parts of) an employee's mobile device through some mobile device management solution. A UE being connected to multiple PDN gateways would allow simultaneous traffic flows from the different domains, including the corporate PDN. For example, the use of a corporate PDN could be limited to certain locations.

Regional public networks and services. Another potential use case could be virtual public networks of a limited scale, for example, connecting users attached to the same set of eNBs. This would set a stage for public safety and social applications such as gaming. The regional boundaries could be of any reasonable level, ranging from households (neighbours) to venues to cities. The PDNs could be either bound to specific areas ("O2Arena", "London"), or universal ("neighbours", "MyCity").

10.5 Road towards implementation

In the previous section, we have described the main design goals of MULTEX, and shown that MULTEX does not require major changes in the network to be implemented. In this section, we discuss the possible PDN discovery mechanisms in existing LTE networks and elaborate on how the discovered PDNs can be used by UEs.

10.5.1 Discovery of available PDNs

The PDNs that a network operator makes available to a particular user, might depend on several factors. First, some users might have static access to certain APNs that is part of the subscriber profile in HSS. Second, the operator might choose to enable certain PDNs only to the users that reside in certain TAs or attached to particular eNBs. Third, PDN discovery and an EPS bearer establishment can be a dynamic process, when the UE detects that certain hosts that can be reached over a PDN that is potentially closer to the user. Finally, the UE has to be aware of and have configuration settings for any APN it wishes to establish.

One of the key limitations with the current LTE networks is that the network cannot trigger the establishment of the default EPS bearers. The network can only trigger the activation of dedicated EPS bearers. However, since every dedicated EPS bearer is linked to a particular active default EPS bearer, the physical path of the flow for the dedicated EPS bearer remains the same as for the linked default EPS bearer.

In the following, we discuss different ways for how UEs can 1) be made aware of the available PDNs, 2) obtain necessary parameters such as APN to establish a default EPS bearer towards the PDN, and 3) discover the services provided and hosts reachable over a particular PDN.

Pre-provisioned APN settings. Depending on the type of the UE (e.g. smartphone, USB dongle, portable router), it might have pre-installed network settings or be provided with software that have the network settings pre-configured. Network settings can also be provisioned over-the-air, either automatically or upon a user request. Network settings typically contain a network profile for the default APN to be used for Internet connectivity, as well as other services the network provides and that require a separate PDN, such as IMS. Users can also edit the existing network settings or create a profile for a new APN to be used for certain services, such as tethering.

We emphasise that most operating systems provide only very limited access to the APN configuration. Some operating systems, for example, Apple's iOS, do not allow to specify what type of services a particular APN provides or disallow defining more than one APN [9]. Even when the user can alter the APN type, which is the case for Android OS that defines 10 possible data connection types, available APN types are limited to specific predefined traffic classes the user has no control over.

MULTEX can work on UEs with static configuration of multiple APNs, but due to configuration constraints imposed by most OSes, the benefits of MULTEX would be limited.

Announcements over an existing PDN connection. Since an LTE UE always has an initial default EPS bearer established and IP address assigned, it is possible to announce the available PDNs by sending the information to the UE via TCP or UDP. Such approach, however, requires a special application to be installed on the UE, which would listen on an arbitrary port in the background to receive the settings.

While a similar technique has been proposed before [91], the presence of such application might not be enough to fully enable connectivity to multiple simultaneous EPS bearers if the application does not have sufficient rights to modify network settings and activate additional EPS bearers. It is therefore necessary for the OS to introduce such privilege and let the user to grant it to the application.

Available PDNs as part of signalling. The initial default EPS bearer is established during the attach procedure, while the establishment of all subsequent default EPS bearers can be triggered by the UE at any time. In both cases, the UE sends the *PDN connectivity request* to the MME, which in turn responds with the *Activate default EPS bearer context request*. The former message, in addition to the negotiated QoS and other parameters, includes Protocol Configuration Options (PCO) and TFT [56]. PCO is used to transfer external network protocol options, such as DNS servers, to the UE. TFT defines up to 16 filters by the means of source and destination IP address prefix and port tuples in uplink and downlink directions. TFTs are then used by the UE and the PGW to map the IP flows to the matching EPS bearers, either default or dedicated. Both PCO and TFT can also be present in *Modify EPS bearer context request* that is used to modify an EPS bearer context, for example in order to apply a new QCI. Typically the modification of EPS bearers is initiated by the network, but UEs can also trigger it by initiating a bearer resource allocation or modification procedure.

Both PCO and TFT structures define the ability to pass arbitrary network-

specific parameters of a variable length [7]. More specifically, PCO has the additional parameters list, where each parameter is identified by a container identifier, most of which being reserved for operator-specific use. Similarly, TFT can be supplemented with a list of parameters.

MULTEX can greatly benefit from one or both of the aforementioned elements in order to convey the available PDNs to the UE without introducing any changes to the signalling protocol or creating performance implications. The network can decide when and which available PDNs to announce depending on the UE's location, the deployment of L-PGWs, and the subscriber profile.

UE initiated DNS assisted PDN discovery.

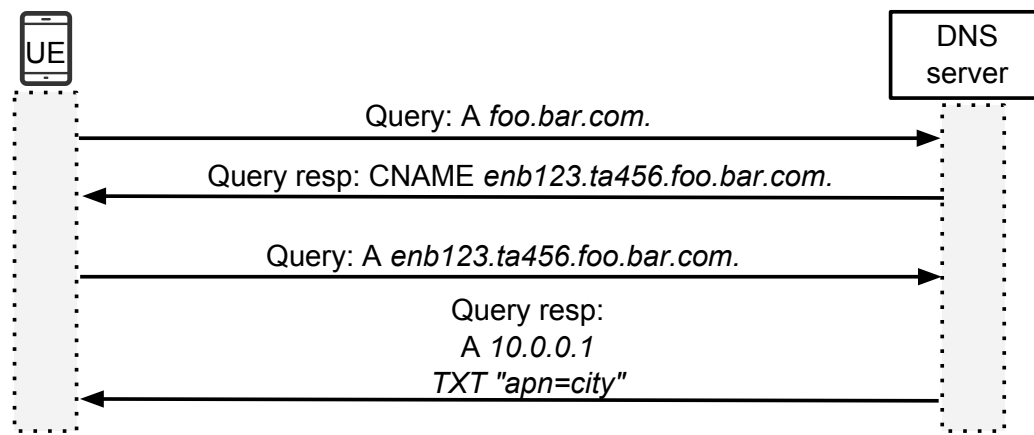


Figure 10.2: DNS based PDN discovery.

All of the aforementioned methods to discover the available PDNs are initiated by the network and meant to be stored by the UE for longer periods. In some cases, it might be beneficial to initiate the PDN discovery from the UE dynamically. The most natural way to realise such discovery is by the means of DNS. In LTE networks, DNS servers used by UEs are controlled by their network operator, thus allowing to manipulate name resolutions in a way that is similar to resolution mechanisms used by CDN [73]. As discussed in 10.2, MME uses contextual attributes, such as TAI or eNB identifier to resolve SGW or PGW for the UE. A DNS server could also use this contextual information to supplement the conventional resolution of users queries for A (AAAA) records with additional pointers to the available PDNs.

In Figure 10.2 we show one example of such resolution process. When the DNS server gets a query from the UE, it responds with an alias (CNAME) that

has an eNB identifier and TAI prefixed to the original host. The UE continues the resolution by sending a query for the alias. As a response, in addition to the actual IP address of the host, it gets information about one or more PDNs that should be preferred for communication with the host. The information about candidate PDNs, such as APN, priority and others, can be conveyed by the means of TXT or NAPTR DNS records.

10.5.2 Routing and decision models

Once the UE is connected to more than one default EPS bearer simultaneously, it needs to decide which route (corresponding to one default EPS bearer or one PDN) IP packets should be sent over. In some cases, one IP destination might be reachable over one and only one PDN, whereas in some other cases the UE could choose between two or more PDNs. If the latter is the case, the UE could use the connectivity to multiple PDNs in order to realise fail-over mechanisms.

In the following, we discuss different routing strategies the UE could use, and elaborate on several routing decision models.

Static TFT-based routing. Depending on the IP protocol version, a single TFT can accommodate up to eight filters to be used for uplink and/or downlink direction. A filter may consist of source or destination IPv4 or IPv6 address prefix, local or remote port range or type and some other parameters. Considering the uplink traffic only, the network can communicate eight IP subnets that are supposed to be reached via a specific PDN.

It is also possible that one IP prefix is present in the TFT of two distinct default EPS bearers. In such case, the UE can selectively decide which PDN to use for outgoing IP packets towards the host present in the filter depending on other criteria. This can for example be the precedence of the filter in the two TFTs.

Dynamic DNS-based routing. The UE initiated DNS assisted PDN discovery described before could be extended to also realise the routing. Once the UE discovers that a host can be reachable over a certain PDN, it can initiate the establishment of the default EPS bearer (if it is not yet established), followed by sending IP packets towards that host over it.

The TXT or NAPTR entry, in addition to the APN, could contain additional parameters, such as whether the UE should terminate the EPS bearer once it finishes the communication with the host. The DNS server can also return multiple TXT or NAPTR entries for a single host in case there is more than one available PDN to be used for accessing the host.

Hints based routing. While the aforementioned TFT and DNS based routing mechanisms are simple, they may not be applicable or flexible enough in all situations. TFT allows store only a limited number of addresses, while DNS-based model might not work for applications using IP addresses instead of host names.

In such scenarios, it might be beneficial to abstract the host entities reachable over a particular PDN to one or more hints that can be expressed as strings. Rather than maintaining the list of domains or IP prefixes, the available PDNs could have hints associated with them. The hints could range from geographical predicates (e.g. "London", "neighbourhood", "mycompany") to the expressions that are specific to certain content types or providers (e.g. "live", "netflix").

Once the UE receives the available PDNs along with the associated hints from the network, it could store them internally and expose to the applications. In turn, the applications demanding content that is potentially available closer to the user, should maintain the list of preferred hints that express such demand. Before initiating a connection towards the host, the application should check whether the most preferred hint is among the available ones stored at the UE. If this is the case, the application would trigger the establishment of a corresponding default EPS bearer (if not yet established) and proceed with sending IP packets over it. If not, the next most preferred hint should be checked and the process continues until no hint can be satisfied and the packets are sent over the initial default EPS bearer.

10.6 A proof of concept of MULTEX

In this section, we demonstrate a conceptual proof of concept of MULTEX that is capable of handling some of the use cases described in Section 10.2. For the discovery of available PDNs, we choose the PCO based scheme, and we use the hints-based mechanism to make routing decisions. We show how MULTEX works on UEs running Android, but similar mechanisms can be implemented on other platforms.

10.6.1 Scenario

We build the proof of concept for the two following scenarios: content delivery from caches within the TA, and peer-to-peer connectivity among the users attached to the same eNB. We therefore rely on the presence of the two additional PDNs, namely "cdn" and "peers", that are announced by a network operator. Our target UE is a smartphone with Android OS that has an active LTE subscription and is

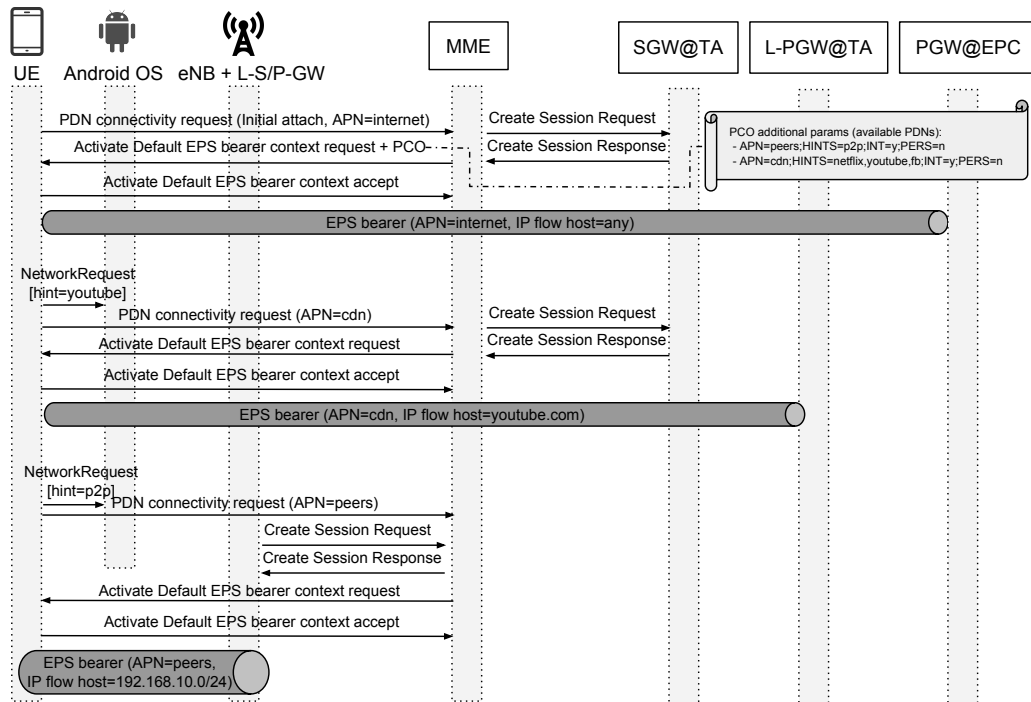


Figure 10.3: A proof of concept of MULTEX.

provisioned with network settings for the default EPS bearer with the APN "internet".

Any application for which a content provider has made caches available within the TA, could benefit from the "cdn" PDN. Examples of such applications could be YouTube, Netflix or Facebook. The "peers" PDN would enable new types of services and applications with peer-to-peer nature, ranging from nearby peers discovery and communication, multimedia streaming, and others.

The sequence diagram for the scenario when the UE becomes attached to the two PDNs is shown in Figure 10.3. In the remainder of this section, we discuss each step in detail.

10.6.2 Local gateways and PDNs

We assume that in addition to a centralised PGW located in EPC and terminating the "internet" APN, the network operator runs two additional L-PGWs. The first L-PGW is assumed to be deployed at the eNB the UE will be attaching to and collocated with the L-SGW. This L-PGW terminates the peer-to-peer PDN, iden-

tified by the APN "peers". The second L-PGW resides at the TA the eNB belongs to and is connected to the SGW serving that TA. This L-PGW hosts the CDN PDN, identified by the APN "cdn".

10.6.3 LTE attach procedure

The UE initiates the LTE attach procedure by sending the *initialUEMessage* that contains the *Attach request* combined with the *PDN connectivity request*. The *initialUEMessage* includes the UE's location information, such as TAI and eNB identifier. The *initialUEMessage* is forwarded by the eNB to the MME, followed by the authentication and signalling setup procedures. The MME accepts the initial attach request by sending the *Attach Accept* and requests the UE to activate the default EPS bearer by sending the *Activate default EPS bearer context request* along with a Non-access stratum (NAS) message. The NAS message, among others, contains the TAI list that the UE does need to send TA updates, assigned QCI, APN, IP address and PCO. The PCO, discussed in Section 10.5, is based on the requested configuration by the UE in the *Attach Request* and typically carries DNS server information.

10.6.4 PDN announcement inside the PCO

Since the MME gets both the UE and the serving eNB contexts as well as the user profile from HSS, it can make informed decisions about which PDNs the UE can be offered with. We assume that the network wishes to announce the availability of both PDNs. To this end, the MME constructs a PCO additional parameters list carrying the information about the additional PDNs and appends it to the PCO structure.

10.6.5 Additional PDNs parameters

Each PDN announced by the MME consists of a number of parameters associated with the PDN. The list could be extended with other parameters, such as authentication parameters or IP protocol version, depending on the needs. In Figure 10.3, the four parameters are used:

- **APN** A globally unique APN that is mapped to one PDN.
- **HINTS** A list of comma-separated hints that will be exposed to applications for routing.

- **INTERNET** A binary flag specifying whether the PDN, in addition to specialised services, also provides Internet connectivity.
- **PERSISTENCE** A binary flag indicating whether the EPS bearer should be terminated upon terminating the TCP or UDP sessions associated with it.

10.6.6 Additional PDN discovery and handling by the UE

On Android, all signalling with the network is handled by the RIL. The establishment of the initial and additional EPS bearer is done by the OS invoking the *RIL_REQUEST_SETUP_DATA_CALL* operation [97]. As a response, the *RIL_Data_Call_Response* structure carrying network parameters, is returned. We propose extending the structure with an additional parameter that is a list containing space-delimited list of available PDNs described previously.

Once the Android OS gets the information about available PDNs from RIL, it needs to implicitly or explicitly expose them to the applications. In the latest Android API, the preferred way of requesting a network connection is through the *requestNetwork* method that is part of the *NetworkRequest* class. Instances of the latter are created using helper classes that include methods for adding or removing network capability and transport type and ability to provide a bearer-specific network specifier. We find the latter suitable for MULTEX to carry the hints of available PDNS.

The applications would perform a lookup of the PDN by specifying the desired hints as the network specifier. The underlying API functions as well as the RIL layer would then check if the hints requested by the application are among the ones it has received from RIL. If such PDN is found but not yet established, *RIL_REQUEST_SETUP_DATA_CALL* with the respective APN would be invoked. Further, if the PDN has the *PERSISTENCE* parameter set, the default EPS bearer associated with the PDN will be terminated by invoking *RIL_REQUEST_DEACTIVATE_DATA_CALL*, once the application is finished with using the network request.

10.7 Evaluation

In this section, we evaluate MULTEX by assessing the support of multiple default EPS bearers in the current LTE networks and devices and discussing the trade-offs and implications that MULTEX would impose on the EPS components.

10.7.1 Support for multiple default EPS bearers in current LTE networks and UEs

UE	Net. interfaces	Level of support
Sierra Wireless MC7304	2xQMI ¹ , 1xPPP	Full support of 3 default EPS bearers
Huawei E392-u12	1xQMI ¹ , 1xPPP	Only one network interface can be active at the same time
Samsung Galaxy Note 4	8xRmNet ¹	Establishment of EPS bearers is controlled by the OS

Table 10.1: Support for multiple default EPS bearers by different UEs.

We conducted a survey on several LTE networks and UEs to check the support for multiple simultaneous EPS bearers. Our test included two LTE commercial networks, one mini PCI express card, one LTE USB dongle, and one smartphone. We used standard mobile broadband subscriptions, which by default had access to several APNs, differentiated by the allocation of public/private IP address to an EPS bearer, firewall restrictions imposed (e.g., all unrelated incoming traffic is blocked or not) and IPv6 support (alone or dual-stacked with IPv4).

Both LTE networks we tested supported at least three multiple default EPS bearers associated by distinct APNs. Each EPS bearer was allocated with a unique IP address, had connectivity to the Internet and, in case of the APN with public IP addresses scheme and no firewall restrictions, was accessible from outside. It was also possible to specify the APN to be used for the initial default EPS bearer that is created during the attach procedure.

The support for multiple simultaneous EPS bearers differed between the three UEs we tested. Table 10.1 shows the support of multiple default EPS bearers for the UEs we tested. While all the UEs allowed us to successfully activate several default EPS bearers, only with the MC7304 UE it was possible to use more than one default EPS bearer simultaneously.

Our survey shows that the support for multiple default EPS bearers is already prevalent in most LTE networks and UEs. We foresee the number of UEs that do not support multiple default EPS bearers to become negligible once VoLTE, which mandates an additional default EPS bearer, gets wider adoption. To this end, the key component that would enable MULTEX and is currently lacking support, is the ability for the UE to get the available PDNs from the network and trigger the activation based on the user or application preference.

10.7.2 Performance of multiple default EPS bearers in real LTE networks

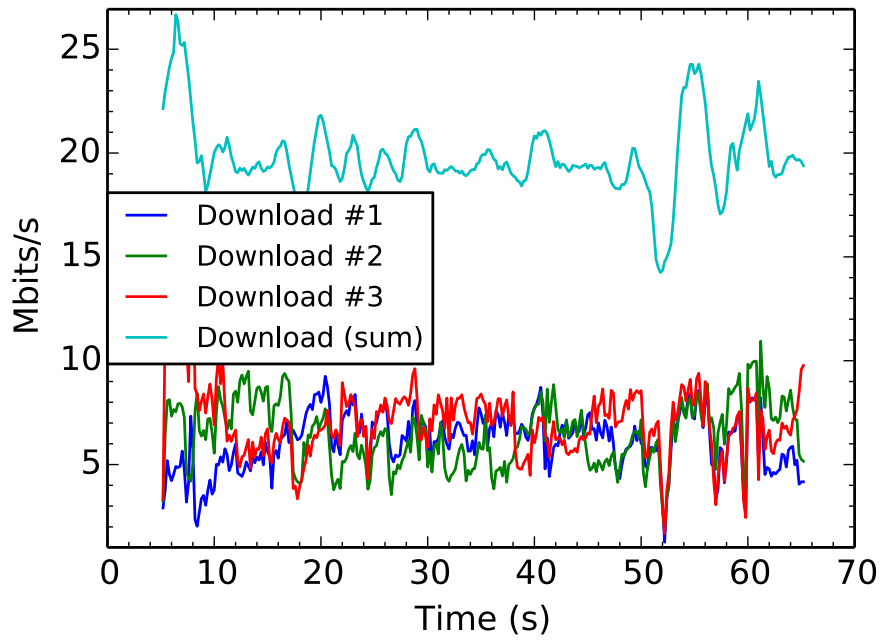
To compare the performance of a single and multiple default EPS bearers, we conducted a simple experiment in a commercial LTE network in Europe, using a standard subscription and Sierra Wireless MC7304 UE. For the multiple default EPS bearers scenario, we established three EPS bearers using QMI¹ interface and PPP interfaces. The measurement consisted of three TCP downloads using connections established over a single or three different EPS bearers. We make two observations about the measurement results provided in Figure 10.4. First, as expected, the aggregate TCP download throughput when using a single EPS bearer is similar to the aggregate throughput when three simultaneous EPS bearers were in use. In both cases, the combined throughput was around 20 Mb/s, which suggests that the use of additional default EPS bearers does not affect the overall throughput. Second, the TCP connections established over three different default EPS bearers highly correlate with each other. This suggests that the variance in throughput is the effect of the scheduler at the eNB that treats individual EPS bearers equally. To further verify this, we repeated our experiment with different TCP queuing disciplines and we observed the same correlation.

10.7.3 The cost of maintaining simultaneous EPS bearers

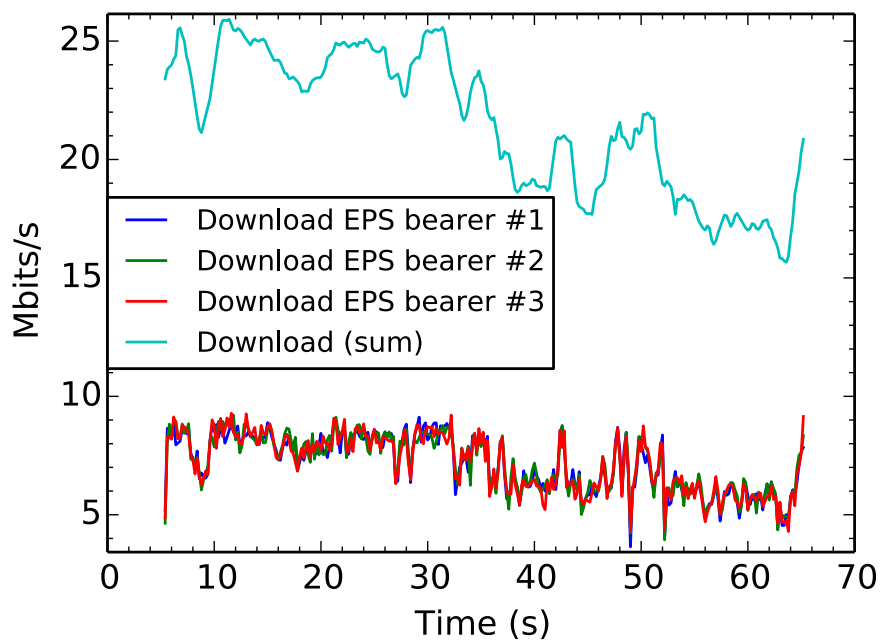
A UE, once attached to a network, exchanges signalling messages with the network to update its state information. Both SGW, PGW and MME stores the TAI of the last TAU sent by a UE, whereas HSS stores the MME that is serving a UE. In turn, the presence of each default or dedicated EPS bearer creates state on different EPS components.

E-RAB establishment. As MULTEX relies on the presence of multiple default EPS bearers, the number of required DRB, S1, S5 and S11 bearers and their respective signalling connections will be proportional to the number of PDN connections established by a UE. While S5 and S11 bearers are established only once and remain active until the EPS bearer is terminated, E-RABs (DRB and S1 bearers) are established or released depending on inactivity timers. When a UE is inactive for certain period of time, the eNB sends the *RRC Connection release message* to

¹Qualcomm MSM Interface (QMI) is a Qualcomm-proprietary interface between Mobile Station Modem (MSM) and attached Terminal Equipment (TE). RmNet interface is a logical interface in QMI framework for data services.



(a) Single default EPS bearer.



(b) Three default EPS bearers.

Figure 10.4: Three parallel TCP downloads over a single (a) and three parallel EPS bearers (b).

a UE to release all active E-RABs. Resuming a suspended E-RAB upon the transition from *IDLE* to *CONNECTED* RRC state involves signalling between MME, SGW, eNB and UE.

Trade-offs of multiple default EPS bearers. While there is a certain cost associated with maintaining each default EPS bearer, the number of resources and signalling required is comparable. First, the network monitors the activity of a UE and imposes inactivity timers regardless of the number of active EPS bearers a UE has [56]. Second, the list of E-RABs to be re-established are always contained in a single signalling message. Third, two physical radio channels for uplink and downlink traffic are shared by all E-RABs, hence any additional default EPS bearers will not require the establishment of additional physical radio channels. Based on this, MULTEX does not create any noticeable overhead in terms of signalling and physical radio resources. We also believe that the cost in terms of extra memory to maintain a state for additional default EPS bearers should be relatively small compared to signalling and traffic off-loaded from EPC and terminated closer to the edge.

Conditional EPS bearer re-establishment. While all DRBs are released upon releasing an RRC connection, 3GPP specifications does not explicitly state whether all DRBs that were previously released should be re-established upon re-establishment of an RRC connection. In other words, the situation when a UE has an RRC connection and some inactive DRBs remains undefined. This might occur when the list of DRBs to activate in the *RRC Connection Reconfiguration* message does not contain all of the DRBs that were previously active. While a UE can request re-establishment of the specific DRBs by sending the *NAS Service Request*, it is meant to re-establish the RRC connection and not individual DRBs. We believe the scenario where only the EPS bearer that triggered data activity gets re-established can be beneficial in terms of reduced signalling and state changes at the respective tunnel endpoints.

10.7.4 Charging, security and mobility.

While the presence of L-PGWs in an LTE network can offload significant portions of traffic from the EPC, the decentralised architecture poses additional challenges for charging, security and management functions. In the following, we discuss some of them in more detail.

Charging. Mobile network operators have several data pricing plans and schemes targeted for different customer segments. These schemes include real-time and

offline charging and both when users are in the home network or roaming. Some or all charging mechanisms supported by a central PGW must also be supported by a L-PGW. To this end, the L-PGW, while terminating portions of traffic closer to the edge, needs to maintain a connection towards EPC to realise the charging function, which in some cases might not be feasible. Operators should therefore try to minimise signalling between L-PGW and EPC, for example by allocating certain data quotas to users periodically and thus avoid the need for real-time connection. Alternatively, operators might consider deploying charging functions at the edge and ensuring that they are in sync with the central databases.

Security. Another major concern in a decentralised environment is to ensure that the level of security is not reduced compared to a centralised architecture. Different PDNs are subject to different security policies and threats, which can lead to different abuses of specialised PDNs, such as IMS [52]. Deploying all security features at or close to every L-PGW might be costly and impractical. One solution to the problem could be security function virtualisation, where one or more security features are realised by a virtual machine residing at the same physical hardware used by L-PGW.

Mobility. Mobility management is one of the central elements of any mobile network. The network has to ensure continuous service availability under vertical or horizontal handovers. However, for PDNs available in a small geographical area, session continuity is not relevant and an EPS bearer can be terminated once a UE moves away. Second, handover procedures are well defined and include scenarios for changing SGW or PGW, that MULTEX could make use of. Finally, it is up to the network to decide which IP addressing schemes to use for different PDNs and users.

10.8 Discussion and conclusions

In this chapter, we presented MULTEX – an approach that transforms LTE UEs into a multi-homed devices with simultaneous access to multiple PDNs with diverse geographical boundaries. We showed that MULTEX is readily deployable in current LTE networks with no modifications to the existing signalling protocols. We provided guidelines for PDN discovery and routing decision mechanisms, and showed the conceptual implementation of MULTEX, realising a few concrete use cases. We also evaluated trade-offs in the implementation and showed that MULTEX does not create any extra overhead in terms of radio resources and signalling.

We believe MULTEX is a natural step towards a next generation, more decentralised and IP-like mobile network. The utility of MULTEX is limited by certain constraints imposed by architectural and technical aspects of LTE. In the remainder of this section, we discuss some of those limitations as well as argue about some of the design aspects of MULTEX.

MULTEX relies on features that are part of LTE standards and already partially or fully supported by LTE equipment. More specifically, the support of multiple default EPS bearers is mandatory for VoLTE capable UEs. MULTEX is the first solution for a UE-based and network-assisted discovery and usage of multiple simultaneous default EPS bearers pointing to geographically distributed L-PGWs that enables traffic offload from EPC.

MULTEX is only applicable to users that are using packet data service from an LTE operator. MULTEX is not applicable for WiFi networks. Still, taking into account the forecasted growth of mobile data traffic as well as increase in IoT devices, we believe that MULTEX is key to enabling more seamless communication for a significant fraction of users and use cases.

SIPTO is conceptually similar to MULTEX and is standardised by 3GPP specifications. SIPTO defines two modes of operation for selected IP traffic offloading: with a single PDN connection and with multiple PDN connections. However, the realisation of SIPTO introduces changes in the signalling protocol. Further, the mode with multiple PDN connections is not well defined as at the time SIPTO specification was created there were very few UEs supporting multiple default EPS bearers. We believe that compared to SIPTO, MULTEX is a self-contained model that does not introduce any changes in the protocols and enables many more use cases, where SIPTO primarily focuses on local (i.e., corporate) versus non-local access.

Part IV

Epilogue

Chapter 11

Conclusions

This thesis has presented a framework for measuring reliability in MBB networks and conducted a large-scale measurement study of Norwegian MBB networks using NNE – a dedicated measurement platform.

The proposed framework uses end-to-end measurements to assess reliability as experienced by the end user. The framework measures reliability at several levels, from the stability of the network connection to the reliability of the data plane and application layer performance. We believe that this framework gives a good basis for describing the overall reliability of a MBB network.

The study uses the proposed framework to measure the reliability in five Norwegian MBB networks for stationary and mobile connections. We used the framework proposed in Chapter 3 and focused on a few selected metrics at each level in the framework. The measurements presented here have demonstrated that there are clear differences in reliability between operators, and that these can be identified and characterised by end-to-end measurements. Networks vary in the stability of connections, in packet loss and delay patterns, and in their ability to support popular applications.

For the stationary scenario, we find that most loss is a direct result of RRC state transitions, both regular and pathological, while the remaining loss is mostly due to activity in the CN. Loss exhibits diurnal patterns and is related to performance degradation episodes that simultaneously impact a significant fraction of geographically-diverse connections. Our study of excessive delays highlighted the complex interplay between physical and transport layer protocols, and showed that the majority of excessive delay events are caused by buffering on both the uplink and the downlink. For the mobile scenario, we established that loss in MBB networks is significant under mobility, and much higher than in the stationary case.

In particular, we find disturbances or handovers between different RATs is a main cause of loss, accounting for about 70% of the total. We further observed that extreme delay episodes lasting several seconds to be more common when nodes are in motion.

We have shown how end-to-end measurements can give insights into the performance of cellular network internals and be used to identify failures and performance problems that are not necessarily captured by the operators' monitoring systems. Our results motivated one of the operators measured to re-examine its network configuration to mitigate loss caused by state transitions. We further showed that using two MBB connections from distinct operators in parallel can potentially give 99.999% availability. These results build the case for independent end-to-end measurement infrastructures like NNE, which allows correlating measurements at different levels and spot potential problems in MBB networks.

As a step towards a next generation, more decentralised, reliable and flexible and IP-like mobile networks, we presented MULTEX – an approach that transforms LTE UEs into a multi-homed devices with simultaneous access to multiple PDNs with diverse geographical boundaries. We showed that MULTEX is readily deployable in current LTE networks with no modifications to the existing signalling protocols.

Acronyms

2G second generation. vii, 2, 3, 15, 24, 89, 110–114, 116, 136, 139–142, 146

3G third generation. vii, 2–4, 7, 21, 22, 24, 43–46, 55, 89, 92, 110–114, 116, 117, 119–121, 134–137, 139, 140, 142, 143, 145, 146, 148–153, 162

3GPP The 3rd Generation Partnership Project. 20–23, 27, 95, 168, 169, 190, 192

4G fourth generation. vii, 7, 22, 24, 55, 134

5G fifth generation. vii, 15, 23, 24, 168

8PSK 8 Phase Shift Keying. 20

AM Acknowledged Mode. 121

API Application Programming Interface. 40

APN Access Point Name. 28, 169, 171, 177–180, 182, 184–187

ARQ Automatic Repeat reQuest. 120, 164

AuC Authentication Centre. 29

BBU Baseband Unit. 24

BSC Base Station Controller. 17, 21, 23, 26–29

BSS Base Station Subsystem. 16, 17, 19–21

BSSGP Base Station System GPRS Protocol. 20, 23

BTS Base Transceiver Station. 17, 21, 26

CA Carrier Aggregation. 23

CDF Cumulative Distribution Function. 1, 81, 82, 85, 87, 93, 98, 100, 102, 103, 108, 114, 116, 151

CDMA Code Division Multiple Access. 15, 22, 56, 63

CDMA2000 Code Division Multiple Access. 21, 24, 71, 72, 77, 78

CDN Content Delivery Network. 178, 181, 185

CID Cell Identifier. 19, 26, 50, 51, 55, 101, 139, 147, 150

CN Core Network. vii, 11, 16, 17, 21, 23, 24, 26, 28, 29, 33, 34, 43, 44, 50, 67, 72, 77, 80, 90–92, 106, 108, 110, 129, 130, 140, 154, 159, 167, 170, 195

CS Circuit-Switched. 16, 18, 26, 28, 29

CSD Circuit Switched Data. 18

CSFB Circuit Switched FallBack. 23

DC Dual Carrier. 22

DL Downlink. 121

DNS Domain Name System. 4, 41, 172, 178, 180–183, 185

DRB Data Radio Bearer. 169, 188, 190

DRX Discontinuous Reception. 27

E-RAB E-UTRAN Radio Access Bearer. 169, 188, 190

E-UTRAN Evolved UTRAN. 24, 26

E_c/I_0 signal-to-noise ratio. 1, 3, 58, 81, 83–85, 89, 119, 126, 135

EDGE Enhanced Data rates for GSM Evolution. 20

eNB E-UTRAN Node B. 23, 24, 26–28, 169–172, 177–179, 181–185, 188, 190

EPC Evolved Packet Core. 23, 24, 26, 43, 168–172, 177, 178, 184, 190–192

EPS Evolved Packet System. vii, x, xi, 4, 5, 28, 33, 34, 47, 58, 67, 77, 134, 135, 142, 169, 170, 172, 173, 176, 177, 179, 180, 182–192

ETSI European Telecommunications Standards Institute. 15, 24, 172

EUL Enhanced Uplink Channel. 121

EV-DO Evolution-Data Optimized. 56

FCC Federal Communications Commission. 36

FDD Frequency Division Duplex. 21–23

FDMA Frequency Division Multiple Access. 19, 22

FTP File Transfer Protocol. 36

GGSN Gateway GPRS Service Node. 18–20, 28, 29, 34, 44, 77, 101, 111

GMSK Gaussian Minimum Shift Keying. 17, 18, 20

GPRS General Packet Radio Service. 1, 18–20, 29, 58

GPS Global Positioning System. 63, 132, 136

GSM Global Standard for Mobile Communications. 1, 15, 16, 18, 19, 21, 23–29, 55, 58, 71, 72

GTP General Packet Radio Service Tunnelling Protocol. 20, 23, 28, 29, 170

HeNB Home eNB. 169

HLR Home Location Register. 17, 29

HSCSD High Speed Circuit Switched Data. 18

HSDPA Hight Speed Downlink Packet Access. 22, 120, 121

HSPA Hight Speed Packet Access. 22, 27, 46, 111, 164

HSPA+ Evolved Hight Speed Packet Access. 22, 46, 55, 111, 164

HSS Home Subscriber Server. 29, 179, 185, 188

HSUPA Hight Speed Uplink Packet Access. 22

HTTP HyperText Transfer Protocol. 2, 3, 36, 38–41, 44, 46, 122–124

ICMP Internet Control Message Protocol. 38, 39, 41

IEEE Institute of Electrical and Electronics Engineers. 22, 32

IETF Internet Engineering Task Force. 32

IMS IP Multimedia Subsystem. 21, 29, 169, 177, 179, 191

IoT Internet of Things. 7, 168, 192

IP Internet Protocol. 19, 20, 23, 28, 29, 32, 34, 45, 57, 58, 66, 67, 77, 111, 134, 135, 142, 150, 167, 168

IPC Inter-process communication. 58

IPPM IP Performance Metrics. 32

IS95 Interim Standard 95. 15

L-PGW Local PGW. 169, 178, 181, 184, 185, 190–192

LA Location Area. 17, 26

LAC Location Area Code. 4, 5, 26, 101, 139, 147–150, 154

LAN Local Area Network. 57, 59, 65

LIPA Local IP Access. x, 169, 172

LLC Logical Link Layer. 19, 20

LTE Long Term Evolution. iv, x, xi, 1–4, 8, 11, 22–29, 43–46, 55, 58, 63, 71, 72, 77, 110–114, 116, 121, 132, 134–137, 139, 140, 143, 146, 148–151, 153, 162, 164, 167–173, 175–177, 179–181, 183, 185–188, 190–192, 196

M2M machine-to-machine. 7, 159

MAC Media Access Control. 18, 19, 96

MBB Mobile broadband. iii, viii, 1–3, 7–11, 22–26, 29, 31–47, 49–52, 54, 56, 57, 61, 62, 64–68, 71–75, 77, 83, 87, 88, 92, 93, 98, 101, 103, 105, 106, 108, 110, 111, 121, 122, 124, 125, 129–132, 134, 137, 150, 151, 153, 157, 162, 164, 167, 195, 196

MCS Modulation and Coding Scheme. 20

MEC Mobile Edge Computing. 24

MGW Media Gateway. 21

MIMO Multiple-Input Multiple-Output. 22, 23

MME Mobility Management Entity. 29, 77, 178, 180, 181, 185, 188, 190

MPTCP MultiPath TCP. 52

MS Mobile Station. 19–21, 26–28

MSC Mobile Switching Center. 17–19, 29

MSM Mobile Station Modem. 188

MTBF Mean Time Between Failures. 4, 80, 157–160

MTTR Mean Time To Restore. 4, 80, 81, 85, 158–160

NAS Non-access stratum. 185

NFV Network Function Virtualisation. x, 24, 172, 173

NNE Nornet Edge. iii, 1, 10, 11, 32, 49–68, 72–74, 77, 78, 81, 101–103, 110, 111, 124, 126, 130, 132, 135, 142, 151, 153, 157, 160, 164, 195, 196

NodeB UMTS Node B. 21, 26, 44, 46, 101

NSF National Science Foundation. 45

NSS Network and Switching Subsystem. 16–18, 22

OFDMA Orthogonal Frequency-Division Multiple Access. 22

OS Operating System. 42, 77, 136

OSI Open Systems Interconnection. 16

PBX Private Branch Exchange. 124

PCO Protocol Configuration Options. xi, 180, 181, 183, 185

PCRF Policy and Charging Rules Function. 29

PDCP Packet Data Convergence Protocol. 21

PDF Probability Density Function. 2, 118, 146

PDN Packet Data Network. x, xi, 4, 28, 29, 167–169, 171–173, 175–188, 191, 192, 196

PDP Packet Data Protocol. vii, 20, 28, 33, 34, 47, 58, 67, 77, 78, 80, 83, 92, 93, 111, 127, 134, 135, 142

PDU Protocol Data Unit. 19, 20, 121

PGW Packet Data Network Gateway. 28, 29, 34, 77, 171–173, 180, 181, 184, 188, 191

PLMN Public Land Mobile Network. 17

PPP Point-to-Point Protocol. 57, 66, 142, 188

PPPd Point-to-Point Protocol daemon. 57, 66, 142, 143

PS Packet-Switched. 18, 23, 24, 26, 28, 29

PSC Primary Scrambling Code. 121

PSTN Public Switched Telephone Network. 17

QAM Quadrature amplitude modulation. 22

QCI Quality of Service Class Identifier. 28, 169, 180, 185

QMI Qualcomm MSM Interface. 188

QoE Quality of Experience. 8, 29, 45, 168

QoS Quality of Service. 16, 18, 20, 29, 180

QPSK Quadrature Phase Shift Keying. 21

QxDM Qualcomm eXtensible Diagnostic Monitor Professional. 120, 121

RAB Radio Access Bearer. 26, 27

RAID redundant array of independent. 59

RAN Radio Access Network. vii, 1, 2, 21, 24, 26, 29, 34, 35, 43, 44, 50, 71, 72, 77, 80–82, 86, 88, 90–92, 94–96, 100–108, 110, 126–129, 140, 154, 155, 159, 160, 162, 164, 167, 170, 172

RAT Radio Access Technology. x, 1–4, 10, 23, 26, 34–36, 40, 42, 46, 51, 55, 58, 68, 77, 83–86, 110–114, 116, 120, 122, 134, 136–151, 153, 154, 196

RF Radio Frequency. 17, 19, 20

RIL Radio Interface Layer. 177, 186

RLC Radio Link Control. 18, 19, 96, 121

RNC Radio Network Controller. ix, 2, 11, 21, 23, 26–29, 44–46, 92, 95, 98, 99, 101–106, 108–111, 148

RNS Radio Network Subsystem. 21

RRC Radio Resource Control. 1, 2, 5, 21, 23, 26–28, 35, 39, 43–45, 51, 55, 57, 58, 81, 86, 90, 94–101, 106, 108–111, 113, 114, 116, 120, 122, 129, 130, 151, 154, 157, 169, 188, 190, 195

RRH Remote Radio Head. 24

RSCP Received Signal Code Power. 135

RSSI Received Signal Strength Indicator. 1, 40, 57, 58, 81–85, 119, 120, 127

RTP Real-time Transport Protocol. 124

RTT round-trip time. 2, 4, 33, 38, 40, 42, 44, 46, 61, 110–115, 121, 151–153

SCTP Stream Control Transmission Protocol. 52

SDN Software-Defined Networking. x, 24, 172, 173

SDR Software Defined Radio. 43

SDU Service Data Unit. 96, 108

SGSN Serving GPRS Support Node. 18–20, 29, 34, 44, 77

SGW Serving Gateway. 26, 29, 34, 77, 171, 172, 178, 181, 184, 185, 188, 190, 191

SIM Subscriber Identify Module. 16, 57

SIP Session Initiation Protocol. 122, 124

SIPTO Selected IP Traffic Offloading. x, 169, 171, 172, 192

SMS Short Message Service. 16, 18

SNDCP Sub Network Dependent Convergence Protocol. 20, 21

SNPDU Sub Network Protocol Data Unit. 20

SSH Secure Shell. 51, 57, 59, 60, 65

SSL Secure Socket Layer. 38

TA Tracking Area. 26, 177, 179, 183–185

TAI Tracking Area Identity. 178, 181, 182, 185, 188

TCP Transmission control protocol. 4, 20, 29, 38–41, 43–46, 61, 67, 122, 123, 180, 186, 188, 189

TDD Time Division Duplex. 21–23

TDMA Time Division Multiple Access. 19

TE Terminal Equipment. 188

TFT Traffic Flow Templates. 173, 180–183

TTI Transmission Time Interval. 20, 23

UDP User Datagram Protocol. 20, 33, 38, 39, 41, 45, 46, 66, 67, 86, 94, 96, 99, 108, 115, 120, 180, 186

UE User Equipment. xi, 5, 17, 21, 23, 26–28, 32–34, 37, 42–45, 77, 80, 111, 120, 121, 131, 135, 140, 142, 146, 147, 150, 153, 154, 167–170, 172, 173, 175–188, 190–192, 196

UL Uplink. 121

UMTS Universal Mobile Telecommunications System. 1, 21–29, 44, 55, 63, 71, 72, 77, 78, 80, 96, 123, 126, 132

USB Universal Serial Bus. 1, 38, 44, 45, 54, 55, 57, 65, 66, 77, 78

UTRAN UMTS Terrestrial Radio Access Network. 21, 24, 26, 96, 101

VLR Visitor Location Register. 17

VoIP Voice-over-IP. 3, 8, 29, 34, 122, 124, 125

VoLTE Voice over LTE. 23, 169, 170, 176, 187, 192

VPN Virtual private network. 43, 171

WCDMA Wideband Code Division Multiple Access. 21, 22, 58, 121

WiMAX Worldwide Interoperability for Microwave Access. 22

XML Extensible Markup Language. 61

Bibliography

- [1] ITU-T recommendation G.711. Pulse code modulation (PCM) of voice frequencies, 1988.
- [2] Hendelsesrapport Problemer i Telenors mobilnett 10. juni 2011. Technical report, Norwegian Post and Telecommunications Authority, 2011.
- [3] Brand Vodafone Rotterdam april 2012 - Incidentonderzoek. Technical report, Agentschap Telecom - Ministerie van Economische Zaken, 2012.
- [4] Mobiperf. <http://www.mobiperf.com>, 2014.
- [5] 3GPP. Local IP Access and Selected IP Traffic Offload (Release 10). TR 23.829, 3GPP, 2011.
- [6] 3GPP. Mobile Radio Interface Signalling Layer 3; General Aspects (Release 12). TS 24.007, 3GPP, 2014.
- [7] 3GPP. Mobile radio interface Layer 3 specification; Core network protocols; Stage 3 (Release 12). TS 24.008, 3GPP, 2015.
- [8] J. Alcaraz, F. Cerdan, and J. Garcia-Haro. Optimizing TCP and RLC interaction in the UMTS radio access network. In *IEEE Network*, 2006.
- [9] Apple Developer. Configuration Profile Key Reference. <https://developer.apple.com/enterprise/ConfigurationProfileReference.pdf>. Accessed: 2016-03-15.
- [10] M. Bagnulo, P. Eardley, T. Burbridge, B. Trammell, and R. Winter. Standardizing Large-scale Measurement Platforms. *SIGCOMM Comput. Commun. Rev.*, 43, 2013.

- [11] V. Bajpai and J. Schönwälder. A survey on internet performance measurement platforms and related standardization efforts. *IEEE Communications Surveys & Tutorials*, 17(3):1313–1341, 2015.
- [12] A. Balachandran, V. Aggarwal, E. Halepovic, J. Pang, S. Seshan, S. Venkataraman, and H. Yan. Modeling Web Quality-of-Experience on Cellular Networks. In *Proc. of MobiCom*, 2014.
- [13] D. Baltrūnas, A. Elmokashfi, and A. Kvalbein. Measuring the Reliability of Mobile Broadband Networks. In *Proceedings of the 2014 Conference on Internet Measurement Conference*, pages 45–58. ACM, 2014.
- [14] D. Baltrūnas, A. Elmokashfi, and A. Kvalbein. Dissecting Packet Loss in Mobile Broadband Networks from the Edge. In *INFOCOM*, pages 388–396. IEEE, 2015.
- [15] D. Baltrūnas, A. Elmokashfi, and A. Kvalbein. MULTEX: Multiple PDN Connections in LTE and Beyond for Enhanced Routing and Services. In *Proceedings of the 5th Workshop on All Things Cellular: Operations, Applications and Challenges*, pages 13–18. ACM, 2016.
- [16] D. Baltrūnas, A. Elmokashfi, A. Kvalbein, and Ö. Alay. Investigating Packet Loss in Mobile Broadband Networks under Mobility. In *IFIP Networking Conference (IFIP Networking) and Workshops*, pages 225–233. IEEE, 2016.
- [17] A. Banerjee, J. Cho, E. Eide, J. Duerig, B. Nguyen, R. Ricci, J. Van der Merwe, K. Webb, and G. Wong. Phantomnet: Research infrastructure for mobile networking, cloud computing and software-defined networking. *GetMobile: Mobile Computing and Communications*, 19(2):28–33, 2015.
- [18] A. Basta, W. Kellerer, M. Hoffmann, H. J. Morper, and K. Hoffmann. Applying NFV and SDN to LTE Mobile Core Gateways, the Functions Placement Problem. In *Proc. of AllThingsCellular*, 2014.
- [19] Q. A. Chen, H. Luo, S. Rosen, Z. M. Mao, K. Iyer, J. Hui, K. Sontineni, and K. Lau. QoE Doctor: Diagnosing Mobile App QoE with Automated UI Control and Cross-layer Analysis. In *Proc. of IMC*, 2014.

- [20] S. Chen and J. Zhao. The requirements, challenges, and technologies for 5G of terrestrial mobile telecommunication. *IEEE Communications Magazine*, 2014.
- [21] Y. Chen, N. Duffield, P. Haffner, W. ling Hsu, G. Jacobson, Y. Jin, S. Sen, S. Venkataraman, and Z. li Zhang. Understanding the Complexity of 3G UMTS Network Performance. In *IFIP Networking Conference*, 2013.
- [22] J. Cho, B. Nguyen, A. Banerjee, R. Ricci, J. Van der Merwe, and K. Webb. SMORE: Software-defined Networking Mobile Offloading Architecture. In *Procs. of AllThingsCellular*, 2014.
- [23] S. Chourasia and K. M. Sivalingam. SDN Based Evolved Packet Core Architecture For Efficient User Mobility Support. In *Proc. of NetSoft*, 2015.
- [24] *Cisco visual networking index: Global mobile data traffic forecast update, 2015 - 2020*. Cisco Systems, Inc., February 2016.
- [25] Clemson University. SciWiNet. <http://sciwinet.org/>.
- [26] *Annual report of the communications regulatory authority (RTT) of the republic of Lithuania*. Communications Regulatory Authority of the Republic of Lithuania, 2016.
- [27] Digium. Asterisk. <http://www.asterisk.org/>.
- [28] S. Eido and A. Gravey. *Advances in Communication Networking: 20th EU-NICE/IFIP EG 6.2, 6.6 International Workshop, Rennes, France, September 1-5, 2014, Revised Selected Papers*, chapter How Much LTE Traffic Can Be Offloaded?, pages 48–58. Springer International Publishing, 2014.
- [29] T. ETSI. 101 297 v8. 2.0 (sep. 2001). *Digital Cellular Telecommunications System (Phase 2+)*.
- [30] K. R. Evensen. MULTI Network Manager (MNM). <http://github.com/kristrev/multi>, 2013.
- [31] A. Faggiani, E. Gregori, L. Lenzini, V. Luconi, and A. Vecchio. Smartphone-based crowdsourcing for network monitoring: Opportunities, challenges, and a case study. *IEEE Communications Magazine*, 52(1):106–113, 2014.

- [32] FCC. Measuring Broadband America. <https://www.fcc.gov/general/measuring-broadband-america>.
- [33] A. Ford, C. Raiciu, M. Handley, S. Barre, and J. Iyengar. Architectural guidelines for multipath TCP development. <http://tools.ietf.org/html/rfc6182>, 2011. Internet RFCs, RFC 6182, ISSN 2070-1721.
- [34] Fortune. T-Mobile Just Suffered an National Network Outage. <http://fortune.com/2016/09/22/t-mobile-overnight-network-outage/>. Accessed: 2016-11-03.
- [35] A. Gember, A. Akella, J. Pang, A. Varshavsky, and R. Caceres. Obtaining In-Context Measurements of Cellular Network Performance. In *Proc. of IMC*, 2012.
- [36] U. Goel, A. Miyyapuram, M. P. Wittie, and Q. Yang. MITATE: Mobile internet testbed for application traffic experimentation. In *USENIX Symposium on networked systems design and implementation*, NSDI '13, 2013.
- [37] Google. Public Dataset: Google Mobile Network Measurements. <https://storage.cloud.google.com/speedometer>, 2013.
- [38] A. Gudipati, D. Perry, L. E. Li, and S. Katti. SoftRAN: Software Defined Radio Access Network. In *Proc. of HotSDN*, 2013.
- [39] E. Halepovic, J. Pang, and O. Spatscheck. Can you GET me now?: Estimating the time-to-first-byte of HTTP transactions with passive measurements. In *Proc. of IMC*, 2012.
- [40] T. Halonen, J. Romero, and J. Melero. *GSM, GPRS and EDGE performance: evolution towards 3G/UMTS*. John Wiley & Sons, 2004.
- [41] H. Holma and A. Toskala. *WCDMA for UMTS: HSPA Evolution and LTE*. John Wiley & Sons Ltd., 4th edition, 2007. ISBN 9780470512531.
- [42] H. Holma and A. Toskala. *HSDPA/HSUPA for UMTS: high speed radio access for mobile communications*. John Wiley & Sons, 2007.
- [43] H. Holma and A. Toskala. *LTE for UMTS-OFDMA and SC-FDMA based radio access*. John Wiley & Sons, 2009.

- [44] J. Huang, F. Qian, Y. Guo, Y. Zhou, Q. Xu, Z. M. Mao, S. Sen, and O. Spatscheck. An In-depth Study of LTE: Effect of Network Protocol and Application Behavior on Performance. In *Proc. of SIGCOMM*, 2013.
- [45] ICT-317669 METIS project. Deliverable D1.1. “Scenarios, requirements and KPIs for 5G mobile and wireless system”. <https://www.metis2020.com/documents/deliverables/>, 2013.
- [46] IEEE. P802.16.3 Project: Mobile Broadband Network Performance Measurements. <http://www.ieee802.org/16/mbnmp/index.html>.
- [47] W. Inc. WhatsApp. <https://www.whatsapp.com/>.
- [48] H. Jiang, Y. Wang, K. Lee, and I. Rhee. Tackling Bufferbloat in 3G/4G Networks. In *Proc. of IMC*, 2012.
- [49] X. Jin, L. E. Li, L. Vanbever, and J. Rexford. SoftCell: Scalable and Flexible Cellular Core Network Architecture. In *Proc. of CoNEXT*, 2013.
- [50] C. Johnson. *Radio Access Networks for UMTS: Principles and Practice*. John Wiley & Sons Ltd., 1st edition, 2008. ISBN 9780470724057.
- [51] J. Joyce. Bayes’ theorem. 2003.
- [52] H. Kim, D. Kim, M. Kwon, H. Han, Y. Jang, D. Han, T. Kim, and Y. Kim. Breaking and Fixing VoLTE: Exploiting Hidden Data Channels and Misimplementations. In *Proc. of SIGSAC*, 2015.
- [53] Y.-H. Kim, Y.-H. Han, M. Kim, Y. S. Park, S. J. Moon, J. H. Lee, and D. K. Choi. Distributed PDN gateway support for scalable LTE/EPC networks. In *Proc. of CCNC*, 2014.
- [54] Y.-h. Kim, H.-K. Lim, K.-H. Kim, Y.-H. Han, and J. Youn. *Advances in Computer Science and Ubiquitous Computing: CSA & CUTE*, chapter A Distributed Mobility Support in SDN-Based LTE/EPC Architecture. Springer Singapore, 2015.
- [55] Z. Koradia, G. Mannava, A. Raman, G. Aggarwal, V. Ribeiro, A. Seth, S. Ardon, A. Mahanti, and S. Triukose. First Impressions on the State of Cellular Data Connectivity in India. In *Procs. of ACM DEV-4*, ACM DEV-4 ’13, 2013.

- [56] R. Kreher and K. Gaenger. *LTE Signaling: Troubleshooting and Performance Measurement*. Wiley, 2016. ISBN 9781118725108.
- [57] A. Kvalbein, D. Baltrūnas, K. Evensen, J. Xiang, A. Elmokashfi, and S. Ferlin-Oliveira. The Nor-net Edge Platform for Mobile Broadband Measurements. *Elsevier Computer Networks special issue on Future Internet Testbeds*, 61:88–101, 2014.
- [58] M. Laner, P. Svoboda, E. Hasenleithner, and M. Rupp. Dissecting 3G Uplink Delay by Measuring in an Operational HSPA Network. In *Procs. of PAM, PAM’11*, 2011.
- [59] N. Larson, D. Baltrūnas, A. Kvalbein, A. Dhamdhare, C. KC, and A. Elmokashfi. Investigating Excessive Delays in Mobile Broadband Networks. In *Proceedings of the 5th Workshop on All Things Cellular: Operations, Applications and Challenges*, pages 51–56. ACM, 2015.
- [60] W. Lehr and J. Chapin. Mobile broadband growth, spectrum scarcity and sustainable competition. In *Telecommunications Policy Research Conference, Arlington VA*, 2011.
- [61] W. Lehr, M. Heikkinen, D. D. Clark, and S. Bauer. Assessing broadband reliability: Measurement and policy challenges. In *Research Conference on Communication, Information and Internet Policy (TPRC)*, 2011.
- [62] L. Li, K. Xu, D. Wang, C. Peng, Q. Xiao, and R. Mijumbi. A Measurement Study on TCP Behaviors in HSPA+ Networks on High-speed Rails. In *Proc. of INFOCOM*, 2015.
- [63] M. Trudel-Lapierre. ModemManager. <http://launchpad.net/modemmanager>, 2013.
- [64] Malaysian Wireless. Digi network outage yesterday caused by “short circuit in power system”. <http://www.malaysianwireless.com/2015/11/digi-network-outage-yesterday-caused-by-short-circuit-in-power-system/>. Accessed: 2016-11-03.
- [65] M. Moradi, W. Wu, L. E. Li, and Z. M. Mao. SoftMoW: Recursive and Reconfigurable Cellular WAN Architecture. In *Proc. of CoNEXT, CoNEXT’14*, 2014.

- [66] S. Muckaden. Myspeedtest: Active and passive measurements of cellular data network performance. In *Proc. ISMA*, pages 1–10, 2013.
- [67] S. Nanda, K. Balachandran, and S. Kumar. Adaptation techniques in wireless packet data services. *IEEE Communications Magazine*, 38(1):54–64, 2000.
- [68] A. Nandugudi, A. Maiti, T. Ki, F. Bulut, M. Demirbas, T. Kosar, C. Qiao, S. Y. Ko, and G. Challen. Phonelab: A large programmable smartphone testbed. In *Proceedings of First International Workshop on Sensing and Big Data Mining*, pages 1–6. ACM, 2013.
- [69] NFV. NFV White Paper. https://portal.etsi.org/nfv/nfv_white_paper.pdf. Accessed: 2016-03-01.
- [70] A. Nikraves, D. R. Choffnes, E. Katz-Bassett, Z. M. Mao, and M. Welsh. Mobile Network Performance from User Devices: A Longitudinal, Multi-dimensional Analysis. In *Procs. of PAM*, 2014.
- [71] Ookla. “Ookla SpeedTest Mobile Apps. <http://www.speedtest.net/mobile/>.
- [72] OpenSignal. OpenSignal. <https://opensignal.com>.
- [73] J. Pan, Y. T. Hou, and B. Li. An overview of DNS-based server selections in content distribution networks. *Computer Networks*, 2003.
- [74] P. Perala, A. Barbuzzi, G. Boggia, and K. Pentikousis. Theory and Practice of RRC State Transitions in UMTS Networks. In *IEEE GLOBECOM Workshops*, 2009.
- [75] M. Portoles-Comeras, J. Manges-Bafalluy, A. Krendzel, M. Requena-Esteso, and A. Cabellos-Aparicio. An Evolutionary Path for the Evolved Packet System. *IEEE Communications Magazine*, 2015.
- [76] F. Qian, Z. Wang, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck. Characterizing Radio Resource Allocation for 3G Networks. In *Proc. of IMC*, 2010.
- [77] F. Qian, Z. Wang, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck. Profiling resource usage for mobile applications: A cross-layer approach. In *Procs. of MobiSys*, 2011.

- [78] Qualcomm Inc. Qualcomm Interface Control Documents. <https://www.qualcomm.com/documents/interface-control-documents>, 2015.
- [79] A. Rao, J. Sherry, A. Legout, A. Krishnamurthy, W. Dabbous, and D. Choffnes. Meddle: Middleboxes for increased transparency and control of mobile traffic. In *Proceedings of the 2012 ACM conference on CoNEXT student workshop*, CoNEXT student '12, 2012.
- [80] F. Ricciato, E. Hasenleithner, and P. Romirer-Maierhofer. Traffic analysis at short time-scales: an empirical case study from a 3g cellular network. *IEEE Transactions on Network and Service Management*, 5(1):11–21, 2008.
- [81] S. Rosen, H. Luo, Q. A. Chen, Z. M. Mao, J. Hui, A. Drake, and K. Lau. Discovering Fine-grained RRC State Dynamics and Performance Impacts in Cellular Networks. In *Proc. of Mobicom*, 2014.
- [82] K. Samdanis, T. Taleb, and S. Schmid. Traffic offload enhancements for eutran. *IEEE Communications Surveys Tutorials*, 2012.
- [83] SamKnows Limited. MOBILE BROADBAND WHITEBOX SAM-KNOWS CONFIDENTIAL BRIEFING. <https://www.samknows.com/broadband/uploads/methodology/SQ304-001-EN-Whitebox-Briefing-Mobile-D01.pdf>, 2012.
- [84] *Methodology and technical information relating to the SamKnows™ testing platform - SQ301-002-EN*. SamKnows™, 2012.
- [85] S. Sen, J. Yoon, J. Hare, J. Ormont, and S. Banerjee. Can they hear me now?: A case for a client-assisted approach to monitoring wide-area wireless networks. In *Proc. of IMC*, 2011.
- [86] M. Z. Shafiq, L. Ji, A. X. Liu, and J. Wang. Characterizing and Modeling Internet Traffic Dynamics of Cellular Devices. In *Proc. of SIGMETRICS*, 2011.
- [87] M. Z. Shafiq, L. Ji, A. X. Liu, J. Pang, S. Venkataraman, and J. Wang. A first look at cellular network performance during crowded events. In *Proc. of SIGMETRICS*, 2013.

- [88] J. Sommers and P. Barford. Cell vs. WiFi: On the Performance of Metro Area Mobile Connections. In *Proc. of IMC*, 2012.
- [89] S. Sonntag, J. Manner, and L. Schulte. Netradar-measuring the wireless world. In *Modeling & Optimization in Mobile, Ad Hoc & Wireless Networks (WiOpt), 2013 11th International Symposium on*, pages 29–34. IEEE, 2013.
- [90] R. N. Staff. Ripe atlas: A global internet measurement network.
- [91] M. Starsinic, A. S. I. Mohamed, G. Lu, D. Seed, B. Aghili, C. Wang, S. Palanisamy, and P. Murthy. An IP-based triggering method for LTE MTC devices. In *Proc. of WTS*, 2015.
- [92] J. P. G. Sterbenz, D. Hutchison, E. K. Çetinkaya, A. Jabbar, J. P. Rohrer, M. Schöller, and P. Smith. Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines. *Comput. Netw.*, 54(8), 2010.
- [93] R. Stewart. Stream Control Transmission Protocol (SCTP). <http://tools.ietf.org/html/rfc4960>, 2007. Internet RFCs, RFC 4960, ISSN 2070-1721.
- [94] S. Sundaresan, W. de Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescapè. Broadband Internet performance: A view from the gateway. *SIGCOMM Comput. Commun. Rev.*, 41:134–145, 2011. ISSN 0146-4833.
- [95] S. Triukose, S. Ardon, A. Mahanti, and A. Seth. Geolocating IP addresses in cellular data networks. In *Procs. of PAM*, 2012.
- [96] F. P. Tso, J. Teng, W. Jia, and D. Xuan. Mobility: A Double-Edged Sword for HSPA Networks. In *Proc. of MobiHoc*, 2010.
- [97] u-blox. Android RIL. Production delivery – Source code. Application Note. [https://www.u-blox.com/sites/default/files/AndroidRIL-SourceCode_AppNote_\(UBX-13002041\).pdf](https://www.u-blox.com/sites/default/files/AndroidRIL-SourceCode_AppNote_(UBX-13002041).pdf). Accessed: 2016-03-01.
- [98] Viber-Media. Viber. <http://www.viber.com/>.
- [99] K. Wang, M. Shen, J. Cho, A. Banerjee, J. Van der Merwe, and K. Webb. MobiScud: A Fast Moving Personal Cloud in the Mobile Network. In *Procs. of AllThingsCellular*, 2015.

- [100] Y. Xu, Z. Wang, W. Leong, and B. Leong. An End-to-End Measurement Study of Modern Cellular Data Networks. In *Proc. of PAM*, 2014.