Investigating Packet Loss in Mobile Broadband Networks under Mobility

Džiugas Baltrūnas, Ahmed Elmokashfi, Amund Kvalbein*, Özgü Alay
Simula Research Laboratory, Oslo, Norway
*Nexia, Oslo, Norway

***Abstract*—Mobile broadband (MBB) connections are often exposed to varying network conditions under mobility scenarios, which can result in packet loss and higher end-to-end delays. Such performance degradation in turn can adversely impact the user experience. In this paper, we study packet loss characteristics of MBB networks under mobility using six measurement nodes that are placed on regional and inter-city trains in Norway for a period of seven months. Our findings show that packet loss is significantly higher for mobility scenarios compared to the stationary. In order to understand the cause of packet loss, we investigate Radio Access Technology (RAT) changes, temporary loss of service, and changes in cells and location area codes (LAC). We surprisingly find that almost all periods with RAT changes involve packet loss. We also observe that 70% of the overall loss happens in periods with RAT changes or temporary loss of service. Further, one third of RAT changes involve connection termination. Our findings highlight the importance of radio access network (RAN) planning and configuration, and provide guidelines to alleviate packet loss in MBB networks.**

I. Introduction

Mobile broadband is becoming the primary Internet access method for a large number of people and services. All types of Internet applications (office, games, video, web, cloud) are now accessed over MBB. According to the Cisco VNI Global Mobile Data Traffic Forecast, global mobile traffic grew by 69% in 2014 and is expected to grow almost tenfold by 2019 [7]. This tremendous growth in MBB demand has put understanding and improving its performance high on the agenda of both decision makers and industry. Several governments have launched activities to measure MBB performance [1, 8]. Measuring and understanding MBB performance is, however, a challenging task. The plethora of scenarios an MBB connection typically experiences requires context-specific studies. Further, there is a lack of measurement methodologies and metrics that are tailored specifically for assessing MBB [4].

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 characterizing and understanding loss a non-trivial task [5]. 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.

ISBN 978-3-901882-83-8 © 2016 IFIP

In this work, we perform a longitudinal practical investigation of packet loss under mobility in operational MBB networks using end-to-end measurements. 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 use this data to compare and characterize loss under mobility to stationary scenarios. We further leverage connection state information to identify the underlying causes of loss. Our measurements and analysis give insights into the characteristics and causes of packet loss under mobility. In summary, this work 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.

The rest of this paper is organized as follows. Section II presents the measurement scenario and data. Section III discusses basic statistics of loss under mobility and proposes a classification methodology to link loss to its likely causes. Sections IV and V analyzes loss in periods with and without connection technology changes respectively. We highlight the related work in Sec. VI and conclude in Sec. VII.

II. Scenario and data

A. Measurement Setup

The measurement setup used in this study consists of six measurement nodes placed on six regional and inter-city trains operating in Norway. This paper is based on data collected by these nodes from July 2014 until February 2015. Figure 1 shows the routes covered by these trains, which includes a reasonable mix of urban and rural areas.

This measurement setup is the mobile subset of the NorNet Edge (NNE) [13], which is a country-wide measurement infrastructure that consists of several hundreds nodes for

Fig. 1: Map of the train routes

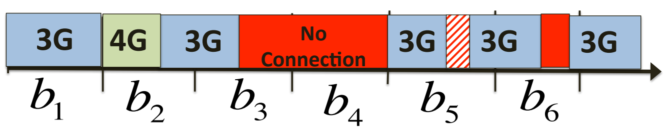

Fig. 2: Typical sequence of connectivity and coverage conditions that MBB connections experience as they move.

measuring the performance and reliability of MBB networks. NNE nodes are single board computers that run a standard Linux distribution and connect to up to four Universal Mobile Telecommunications System (UMTS) operators and one Code-Division Multiple Access (CDMA) 1xEV-DO operator. Like other NNE nodes, train nodes connect to up to four UMTS operators via Huawei E392-u12 modems that support up to Long Term Evolution (LTE) Category 3, and one CDMA 1xEV-DO operator. In this study, however, we limit ourselves to studying two UMTS operators, Telenor and Netcom, because these operators are the only operators that provide LTE service and run their own radio access and core networks. Software running on NNE nodes ensures that MBB connections are always alive, and collects connection state information. In particular, we monitor and record the RAT, which can be *No service*, 2G, 3G or LTE; different signal quality indicators (e.g., RSSI, Ec/Io, and RSRQ); network attachment information (e.g., serving cell identifier, location area code, and tracking area code); and Radio Resource Control (RRC) state. To measure packet loss, we send a 20-byte UDP packet every second over each connection to an echo server that is part of NNE backend and then record a reply packet from the server. A packet is considered lost if we do not receive a reply within one minute. Further, we aggregate the data into five minute bins and calculate loss percentage for each bin. Both measurement data and metadata are periodically transferred to a server and imported into a database. We also 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.

B. *Measurement scenario*

Figure 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 looses its Packet Data Protocol (PDP) context (3G) or Evolved Packet System (EPS) bearer (LTE), which is a tunnel that connects the user equipment (UE) to the core network (CN). Consequently the connection looses 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 neighbor cells, but it looses connectivity to the current cell before completing the handover. Another cause can be failures during inter-RAT handovers, which are known to happen [14]. Finally, connections also experience periods with 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 paper, 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 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 study.

C. *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 artifacts; a different system or hardware may cope better or worse. Next we describe these artifacts 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 reseting the context (bearer). As a result, the operator's firewall drops all incoming packets from these connections, causing 100% packet loss during these periods.

Other artifacts 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 Ec/Io or 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 to 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 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 [5]. 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.

2) 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 > 0. 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.

3) We remove all 5-minute bins that coincide with known server-side maintenance.

4) 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.

5) 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 five minute bins, 38417 from Telenor and 25420 from Netcom.

III. Loss under mobility

In this section, we give a general overview of loss characteristics in mobile networks. We investigate the effect of mobility, and establish the very different loss rates in a mobile vs stationary scenario. We also look at loss in different RATs, before we proceed to classify loss under mobility and relate it to handover events.

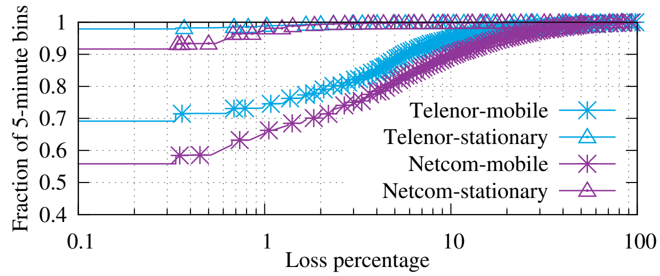

Fig. 3: Overall loss rate and the effect of mobility. Much higher loss rates observed when mobile in both networks.

A. The effect of mobility

Figure 3 shows the overall loss rate for the two measured networks, when the nodes are stationary and moving. It is clear from the figure that loss 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%. This large difference in loss rate between mobile and stationary nodes largely motivates this study. In our previous work [5], we analyzed loss in a stationary scenario. We concluded that loss rates are low in general and that causes are related to misconfiguration of the radio access controller or lie beyond the RAN. In the remainder of this paper, we will dissect and seek to explain what causes loss to be so much higher under mobility. 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 Sec. V). 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 [14].

B. 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 challenged areas (with limited coverage), and that our measurements are therefore not representative for normal 2G behavior. In particular, our dataset contains only 190 and 148 2G bins for Telenor and Netcom, respectively. 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 5 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

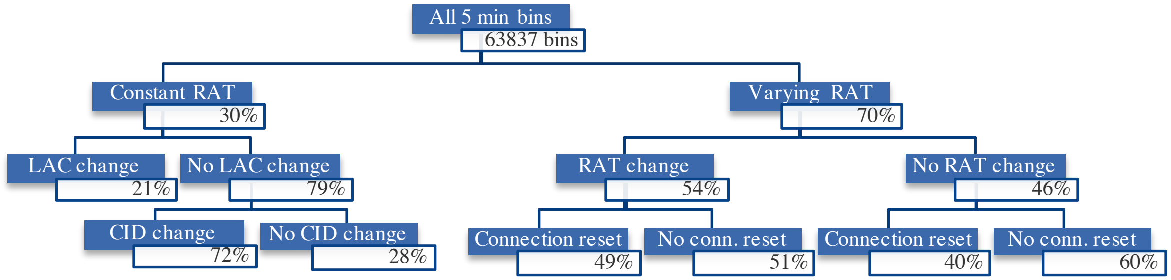

Fig. 4: Classification of loss. The numbers given are percentages of lost packets relative to the parent category. In total, there are 63837 5-minute bins in the dataset. About 229992 packets were lost during these bins.

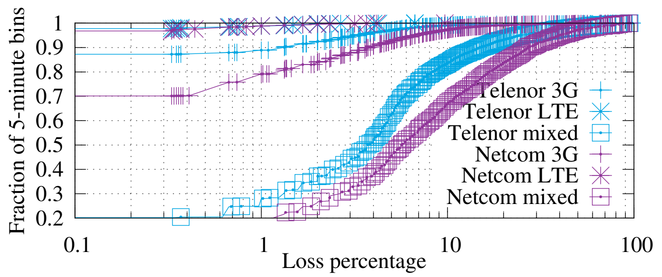

Fig. 5: Loss rate for 3G, LTE and when a RAT change is involved (mixed). Most loss happens in mixed category, LTE performs the best.

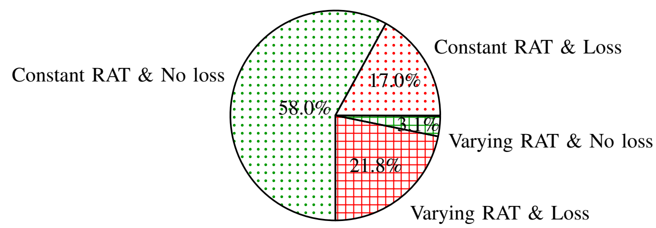

Fig. 6: The percentages of lossy and non-lossy 5-minute bins for constant and varying RATs.

loss, while 14 % (Telenor) to 40 % (Netcom) of 3G-only bins experience loss.

The most striking observation from Fig. 5 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 Sec. IV.

C. Classification of 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 Fig. 4. This classification captures the connectivity states shown in Fig. 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 Sec. II. Inspired by the observation in Fig. 5 that loss is much higher when there is a RAT change, we first split the bins into constant and varying RAT. Constant RAT bins are characterized 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 6 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 behavior is a normal part of a RAT transition, but it also sometimes appears without a resulting RAT change.

Recall from Sec. II 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 Fig. 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 Sec. IV.

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 cell ID (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 Fig. 4 confirms that this is the case: 78% of loss in constant RAT periods happen in bins with either LAC or CID changes. Section V provides a more detailed investigation of causes of loss in constant RAT periods.

IV. VARYING RAT

In this section, we analyze loss in bins with varying RAT, which constitutes 25% of all bins and are responsible for 70% of the overall loss. As mentioned in Sec. III, 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 characterized by one predominant RAT (2G, 3G or LTE), but with the presence of one or more *No service* episodes throughout the bin. Figure 7 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.

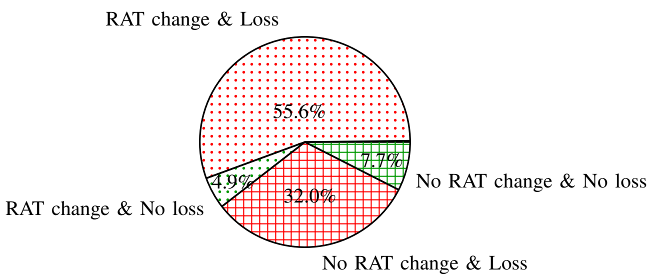

Fig. 7: The percentages of lossy and non-lossy 5-minute bins with varying RATs split by whether more than one RAT is observed.

A. 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 both in 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 8 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

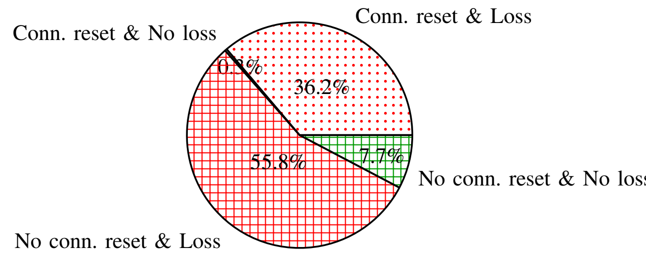

Fig. 8: The percentages of lossy and non-lossy 5-minute bins for varying RAT with a RAT change split by the presence of connection resets.

change and connection resets include packet loss compared to 88% of the bins with RAT changes but 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*.

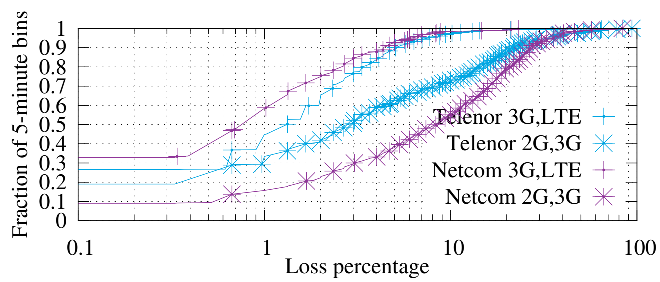

Fig. 9: Involved RATs for the 5-minute bins with a RAT change and no connection resets. Clear differences between networks. More loss during 2G/3G handovers.

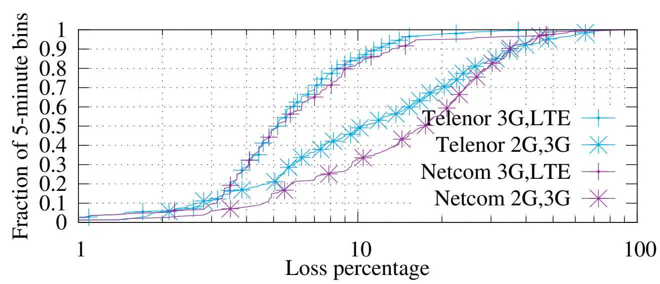

Fig. 10: Involved RATs for the 5-minute bins with a RAT change and one or more connection reset. Much more loss compared scenario without disconnects, minimal differences between networks, least lossy are bins with 3G/LTE handovers.

To further analyze loss related to RAT changes, we identify all distinct RATs present in each 5-minute bin with an inter-RAT handover [1]. Figures 9 and 10 shows the distribution of loss in bins with RAT changes split by the involved RATs. Figure 9 shows bins *without* connection resets, while Fig. 10 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

[1]There 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.

when 2G is among the distinct RATs. Bins where 2G is involved occur in poorly covered areas that are characterized 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 handovers, which initially result in packet loss followed by the connection reset.

To quantify the impact of loss in bins with RAT changes, we count the number of consecutively lost packets in each loss episode. This captures the burstiness of packet loss and we term such consecutively lost packets a loss run. 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 11 shows the probability density function of the loss run size distribution for different RAT combinations for bins with connection resets.

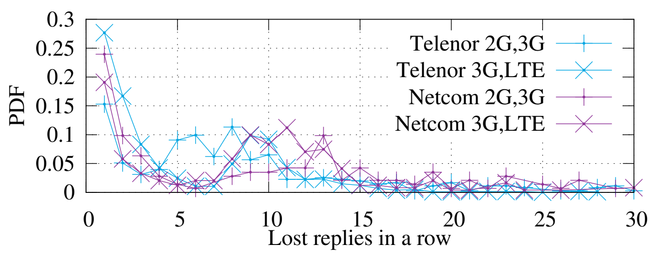

Fig. 11: Loss runs for involved RATs when there is a RAT change and one or more connection reset. Several modes around 10 packets in both networks. Different modes for different sets of involved RATs.

While a sizable 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. These cases can occur when the modem is busy with trying to exchange signaling 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 Fig. 11.

B. 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.

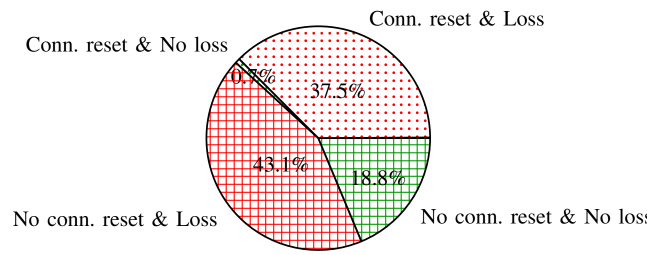

Fig. 12: The percentages of lossy and non-lossy 5-minute bins for varying RATs with no RAT change split by the presence of connection resets.

As for the periods with inter-RAT handovers, some of the bins with no RAT changes have one or more connection resets. Figure 12 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.

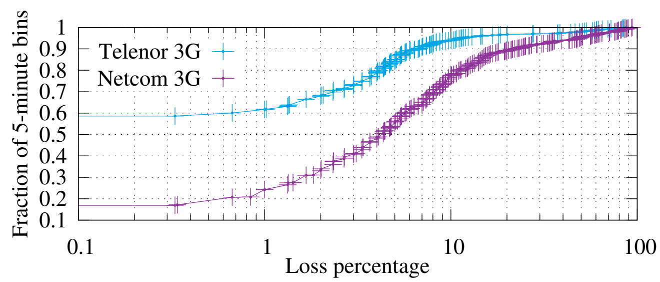

Fig. 13: Individual varying RATs for the 5-minute bins when RAT does not change and there are no connection resets. Less loss in 3G bins with varying RAT in Telenor compared to Netcom.

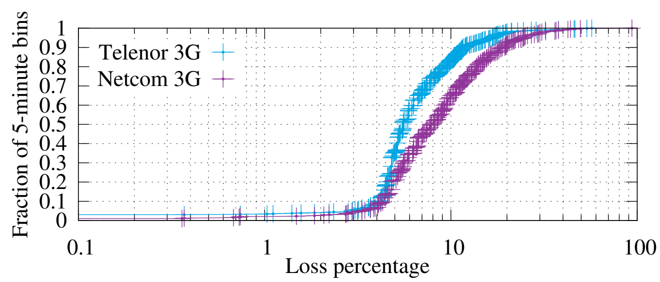

Fig. 14: Individual varying RATs for the 5-minute bins when RAT does not change and there is one or more connection reset. Differences between networks are minimal, but there is slightly less loss in 3G bins with varying RAT in Telenor compared to Netcom.

Next, we look at loss rate distributions for bins with *No service*. Figure 13 and 14 show the loss rate for the bins without or with one or more connection 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 Sec. IV-A.

To quantify the impact of loss in bins with RAT changes, we now look at the distribution of loss runs sizes. Figure 15

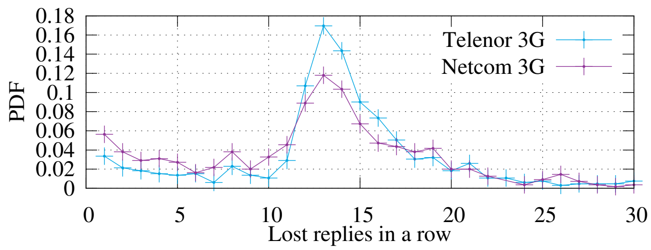

Fig. 15: Loss runs for individual varying RATs when there is no RAT change, but one or more connection reset in a 5-minute bin intersecting with the loss run. All modes in a shorter range from 11 to 16, with a peak at 13 packets for both networks.

shows the 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 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.

V. Constant RAT

This section investigates bins that are characterized by constant RATs (i.e., no inter-RAT handovers), which are located on the left most subtree in Fig. 4. During these periods a connection may experience LAC and cell changes as well as channel quality degradation.

As shown in Sec. III, 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 16a 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

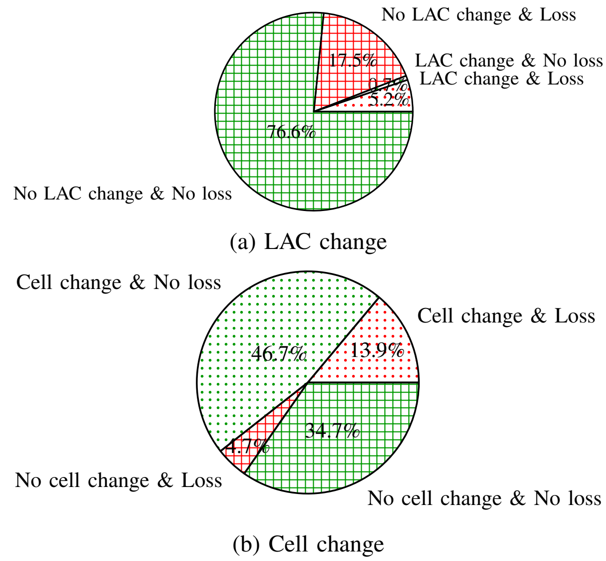

Fig. 16: The percentages of lossy and non-lossy 5-minute bins for constant RATs split by the LAC and cell changes.

completed, in-flight packets will be re-routed to the new cell. Inter-LAC handovers are slightly more challenging, since they involve additional coordination between several RNCs [2], i.e., the handover procedure takes longer to complete compared to cell changes within the same LAC.

Figure 16b 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 17 shows loss rate distribution for bins with LAC changes, cell changes, and no changes when the connections are on 3G or LTE.

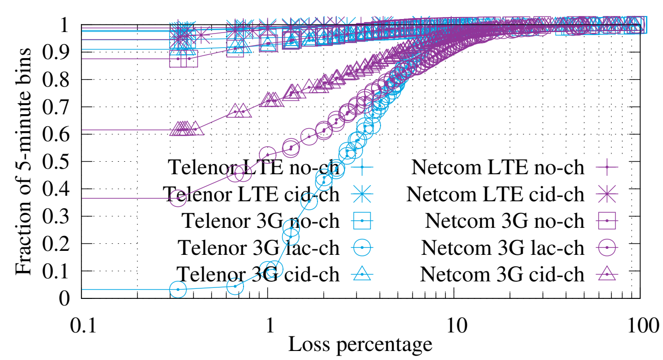

Fig. 17: Loss rate in a 5-minute bin for 3G and LTE RATs split by the presence of one or more LAC change or cell change, or none of them. In both networks the highest loss occurs when there is a LAC change involved. In Telenor, bins with LAC changes are more lossy compared to Netcom, whereas in Netcom 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.

Loss rates are evidently higher in bins with LAC changes with clear differences between operators. Almost all LAC

[2] In 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.

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.
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 looses sight of the old towers during the movement while the procedure is ongoing.

[labelindent=0pt,itemindent=0pt]

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

Table I 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 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.

VI. Related work

There has been a growing interest in performance and reliability measurements of MBB networks. Regulators need measurements to monitor how operators fulfill their obligations, and as a baseline for designing regulatory policies. On the other hand, operators are interested in operational instability and anomalies to identify problems in their networks. There are mainly three approaches for measuring the performance and reliability of MBB networks: (i) crowd-sourced results from a large number of MBB users [2, 15, 19, 23] , (ii) measurements based on network-side data [10, 11, 21, 22] and (iii) measurements collected using dedicated infrastructure [4, 12, 20]. In this paper, 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. The long-term end-to-end measurements lead to a better quality dataset without requiring access to network-side logs, which are typically only available to operators.

Several studies focused on the causes of packet loss in MBB networks. Different groups blamed RRC state transitions [5, 6, 16–18] and showed that the operators do not always configure their RRC state machines according to the standard guidelines leading to significant loss during state demotions. Gember et al. 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 [9]. Xu et al. discussed the effect of bursty packet arrivals and drop-tail policies employed by the operators [25]. RNC-level performance analysis of UMTS networks identified correlations between RTTs and loss and their dependency on diurnal patterns and overloaded NodeBs [6]. One study showed that most transport-layer packet loss is related to physical layer retransmissions and can be reduced by buffering [11]. Another study presented a framework for measuring the user-experienced reliability in MBB networks, and showed how both radio conditions and network configuration play important roles in determining reliability [4]. Both of these studies consider only stationary scenarios, while in this paper we focus on mobility scenarios where signal quality is varying as well as handovers are present.

Packet loss has also been investigated for mobility scenarios. Li et al. [14] 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. Similar observations were made in a study by Balachandran et al. [3], showing that most HTTP sessions with inter-RAT handovers are abandoned. Tso et al. [24] measured HSPA performance on the move to be greatly different from static HSPA performance. In particular, they observed that the final results of handovers are often unpredictable and that UDP packet loss at least doubles during handover periods. Although these studies considered different aspects of packet loss for stationary and mobility scenarios, to the best of our knowledge, there has been no comprehensive study that characterizes packet loss in 3G and LTE networks and compares the mobility and stationary scenarios. Along with the end-to-end measurements used in this work, we further leverage connections' metadata and state information to identify the underlying causes of loss.

VII. Discussion and conclusions

This paper has analyzed the causes of loss in MBB networks under mobility. The observed loss rates are much higher than in the stationary case [5]. In particular, 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.

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, this might not always be sufficient. We therefore acknowledge that network side data or measurements from the RAN could give more insights into the potential causes and assist in improving the network.

VIII. Acknowledgements

This work was partially supported by the European Union's Horizon 2020 research and innovation program under grant agreement No. 644399 (MONROE) and by by the Norwegian Research Council under grants 209954 (Resilient Networks 2) and 208798 (NorNet). The views expressed are solely those of the author(s). We would like to thank NSB for hosting measurement nodes aboard their operational passenger trains.

References

[1] *Annual report of the Communications Regulatory Authority (RTT) of the Republic of Lithuania*, 2012.

[2] Mobiperf. http://www.mobiperf.com, 2014.

[3] 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.

[4] D. Baltrūnas, A. Elmokashfi, and A. Kvalbein. Measuring the Reliability of Mobile Broadband Networks. In *Proc. of IMC*, 2014.

[5] D. Baltrūnas, A. Elmokashfi, and A. Kvalbein. Dissecting Packet Loss in Mobile Broadband Networks from the Edge. In *Proc. of INFOCOM*, 2015.

[6] 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 *Proc. of IFIP Networking*, 2013.

[7] *Cisco visual networking index: Global mobile data traffic forecast update, 2014 - 2019*. Cisco Systems, Inc., February 2015.

[8] FCC. 2013 Measuring Broadband America February Report. Technical report, FCC's Office of Engineering and Technology and Consumer and Governmental Affairs Bureau, 2013.

[9] A. Gember, A. Akella, J. Pang, A. Varshavsky, and R. Caceres. Obtaining In-Context Measurements of Cellular Network Performance. In *Proc. of IMC*, 2012.

[10] 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.

[11] 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.

[12] 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.

[13] A. Kvalbein, D. Baltrūnas, J. Xiang, K. R. Evensen, A. Elmokashfi, and S. Ferlin-Oliveira. The Nornet Edge platform for mobile broadband measurements. *Elsevier Computer Networks special issue on Future Internet Testbeds*, 2014.

[14] 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.

[15] A. Nikravesh, D. R. Choffnes, E. Katz-Bassett, Z. M. Mao, and M. Welsh. Mobile Network Performance from User Devices: A Longitudinal, Multidimensional Analysis. In *Procs. of PAM*, 2014.

[16] P. Perala, A. Barbuzzi, G. Boggia, and K. Pentikousis. Theory and Practice of RRC State Transitions in UMTS Networks. In *IEEE GLOBECOM Workshops*, 2009.

[17] 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.

[18] 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.

[19] J. P. Rula, V. Navda, F. Bustamante, R. Bhagwan, and S. Guha. "No One-Size Fits All": Towards a principled approach for incentives in mobile crowdsourcing. In *Proc. of IMC*, 2014.

[20] 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.

[21] 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.

[22] 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.

[23] J. Sommers and P. Barford. Cell vs. WiFi: On the Performance of Metro Area Mobile Connections. In *Proc. of IMC*, 2012.

[24] F. P. Tso, J. Teng, W. Jia, and D. Xuan. Mobility: A Double-Edged Sword for HSPA Networks. In *Proc. of MobiHoc*, 2010.

[25] 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.