The Good, the Bad and the Implications of Profiling Mobile Broadband Coverage

Andra Lutu*, Yuba Raj Siwakoti, Özgü Alay, Džiugas Baltrūnas, Ahmed Elmokashfi

Simula Research Laboratory, P.O. Box 134, 1325 Lysaker, Norway

Abstract

Pervasive coverage and continuous connectivity of Mobile Broadband (MBB) networks are common goals for regulators and operators. Given the increasing heterogeneity of technologies in the last mile of MBB networks, further support for seamless connectivity across multiple network types relies on understanding the prevalent network coverage profiles that capture different available technologies in an area. Correlating these coverage profiles with network performance metrics is of great importance in order to forestall disturbances for applications running on top of MBB networks. In this paper, we aim to profile MBB coverage and its performance implications from the end-user's perspective along critical transport infrastructures (i.e., railways in Norway). For this, we deploy custom measurement nodes on-board five Norwegian inter-city trains and we collect a unique geo-tagged dataset along the train routes. We then build a **coverage mosaic**, where we divide the routes into segments and analyze the coverage of individual operators in each segment. We propose and evaluate the use of hierarchical clustering to describe prevalent coverage profiles of MBB networks along the train routes and classify each segment accordingly. We further analyze the areas we classify with each profile and assess the packet-loss and HTTP download performance of the networks in those areas.

Keywords: Mobile Broadband, Mobile Coverage, Network Performance Measurements, Data Analysis, Machine Learning

1. Introduction

Mobile Broadband (MBB) access to Internet enables operators to join mobility and communications towards the common goal of offering subscribers performance and efficiency in highly dynamic mobile scenarios. However, Internet access under mobility brings a number of challenges, including high probability of service interruptions. A popular example of such scenarios is the case of travelers regularly commuting on public transport infrastructures, such as inter-city trains. In this context, tens or hundreds of passengers try to access the Internet simultaneously for entertainment, communication and workrelated tasks, all while moving at high speeds. During the last years, railway operators throughout the world have been testing and providing commercial Internet connectivity solutions aimed at enabling on-board Internet services to train passengers. Various types of communication solutions have been advanced [1, 2], including cellular solutions, WLAN-based solutions or hybrid terrestrial/satellite solutions.

The performance of cellular-based solutions for on-board connectivity highly depends on the MBB coverage around the railway lines. MBB operators are the main providers of coverage maps for other stakeholders, including regulators, subscribers or businesses such as public transport operators. These coverage maps usually define the status of one radio access technology (RAT) in a region for a MBB operator. However, they do not offer information on how end-users actually experience the distribution of different RATs in the same geographical region. For example, the existence of 4G connectivity in an area does not mean that all end-users are able to use that technology and it does not necessarily guarantee a good user experience for the end-users that are able to access it. This may be due to a number of factors, including specific geography of the area, variable train speed, number of passengers in the train or congestion in the network. Given the increasing heterogeneity of technologies in the last mile of MBB networks, user experience highly depends on support for seamless handovers across multiple network types. Therefore, identifying the *network coverage profiles* that capture the distribution of all available technologies in the same area from the end-user experience is very important.

In this paper¹, we focus on profiling the MBB coverage along the critical railway infrastructure in Norway from the end-user perspective. In Figure 1, we summarize the workflow we follow in this paper in order to achieve this goal. Our main objective is to build a coverage mosaic, where we classify and characterize railway route segments based on the distribution of RATs along that segment and on how end-users traveling along that route would experience the service. For this, we use a vast dataset that we collect through periodic measurements from custom devices that we strategically place on-board several passenger trains. The dataset is pestered by numerous challenges, including high volume, the mixture of spatio-temporal coordinates and the presence of categorical variables (i.e., the RAT value). Furthermore, depending on the deployment of base stations along the railway routes, the distribution of different RATs highly varies from one segment to another. Some operators rely

^{*}Corresponding author. Tel.: (+47) 458 36 244

Email address: andra@simula.no(Andra Lutu)

¹This paper is an extension of prior work [3] the authors published in the 2016 IFIP International Workshop on Traffic Monitoring and Analysis.



Figure 1: The workflow we follow to generate the coverage mosaic maps.

on thresholds that are imposed on statistical descriptors to characterize coverage and classify different regions in coverage categories. Even when envisioning a simple intuitive classification such as "good" coverage (where we have dominant 3G and 4G) or "bad" coverage (where we have dominant 2G or no service), there is no consensus on what should be the threshold in terms of percentage of 3G/4G presence in a certain area in order to label that area with good coverage. Thus, it is challenging and cumbersome to use statistical descriptors to define good and bad coverage. In this paper, we leverage ideas from machine learning to help us overcome limitations of using classification rules based on statistical descriptors to characterize coverage. More specifically, we evaluate the use of *hierarchical clustering* to characterize the distribution of different RATs from individual MBB providers along the train routes. Clustering facilitates the efficient manipulation of this dataset and enables us to determine the salient coverage profiles, characterize them and then classify the route segments with the proper profile.

The coverage mosaic we produce successfully captures the mixture of available RATs as experienced by the end-user inside the train. Two main coverage profiles emerge from our analysis, one where 3G is dominant (which we further title *"profile A"*) and another where *No Service* is dominant (which we further title *"profile B"*). This validates the intuition within the community regarding *"good"* and *"bad"* coverage profiles, we demonstrate the need for repetitive measurements (at least 5-10 measurement runs) in order to profile the coverage of a certain area.

We finally investigate the implications of the coverage profiles on network performance by analyzing basic QoS metrics and application performance by analyzing HTTP performance metrics. Regarding network performance, we are able to pinpoint the areas with Profile B ("bad") coverage as trouble zones with poor performance. These are areas where operators need to reiterate their coverage evaluation measurements for proper characterization. We further validate these results with HTTP download analysis. Considering that mobile operators often deploy performance enhancing proxies, we carry out our analysis both on the web port (TCP port 80) and a different port (TCP port 85) in order to understand the impact of proxies on the application performance for different profiles. For both operators, we observe high percentages of successful downloads for Profile A, while noting very low percentages of successful downloads for Profile B. We also find evidence of the presence of a web proxy in Telia's network, which slightly increases the HTTP success rate in both profiles.

2. Measurement Setup and Datasets

In this section, we present the measurement infrastructure that we use in this paper, the measurements we deploy and the dataset we collect. We summarize the terminology we use throughout this paper in Table 1.

2.1. Measurement Infrastructure

We use the NorNet Edge [4] (NNE) dedicated mobile broadband measurement platform that is designed to measure the performance and reliability of mobile broadband networks from the user's perspective. NNE nodes are single board computers that run a standard Linux distribution and connect to multiple MBB operators at the same time. The node connects to the different broadband providers via Huawei E392-u12 modems supporting 4G/LTE connectivity. The software running on the NNE nodes ensures that the MBB connections are alive and also collects network connection information. All the data collected on the node are periodically transferred to a server hosted in the back-end and then imported into a database. Note that the data collected using this platform depends on the hardware used and, more specifically, the particular implementation of the device logic.

In order to measure the mobile scenarios, we expand the NNE testbed to include 6 custom NNE measurement devices (i.e., NNE nodes) active on the NSB² regional trains in Norway. Figure 2 shows the routes covered by these trains on the

²NSB is a government-owned railway company operating most passenger trains in Norway.

(NNE) Node	A NorNet Edge node is a small com-		
	puter that is used for MBB measure-		
	ments.		
Route	A fixed geographical path following a		
	train route from one city to another city		
	in Norway.		
Run	A one way trip in a route when the NNE		
	node collects data.		
GPS point	The point in space where a GPS read-		
	ing is taken.		
Metadata	Network context information (e.g.,		
	RAT, signal strength, signal to noise ra-		
	tio, etc.).		
(Geo-tagged)	A data-point tagged with a geographi-		
Data-point	cal location (GPS coordinates).		
Grid block	Square area that results from superim-		
	posing a grid on the map of Norway for		
	geographical binning of the geo-tagged		
	data points.		
(Route) Seg-	Portion of the railway route delimited		
ment	within a grid block.		
Coverage	The distribution of RATs [2G, 3G, 4G,		
Chart	no service] over the geo-tagged data		
	points along a route segment in one		
	drive run.		

Table 1: Terminology.

population density map of Norway. As illustrated in the figure, the routes traverse a reasonable mix of urban and rural areas. The regional trains that host our measurement nodes run periodically on 4 different national *routes*³, covering over 2,500 km. We define each one-way train trip on a certain route as a *run*. We collect data from these nodes for the two largest MBB operators in Norway, Telenor and Telia for 5 months (from November 2014 to March 2015). In Table 2 we summarize the number of runs we have for each operator on each route and different direction. The total number of runs per route ranges from 60 to 150 runs and depends on the public schedule of the trains hosting the nodes.

The measurements we collect from the NNE nodes operating aboard passenger trains in Norway mimic the user experience of the passengers in the train. Under mobility scenarios, as with all user equipment (e.g. smartphones or modems), the modems will report a single RAT value, which is the best available RAT, regardless of how many RATs are available in an area. In other words, our platform behaves in the same way as the user equipment would and this is exactly what we aim to measure with our platform: the RAT experienced by the end user.

2.2. Datasets

In this section, we describe the measurements we deploy on the NNE mobile nodes and the resulting datasets. We transfer all the data we collect on the node to a server we host in the



Figure 2: Spatial locations of the MBB measurements from the NNE nodes operating aboard the NSB trains overlayed on the population density map of Norway. In the population density map, red color indicates highly populated areas such as cities whereas yellow color indicates thinly populated rural areas.

Route	Number of Measurement Runs	
	Telenor	Telia
Oslo - Voss	125	99
Oslo - Stavanger	138	147
Oslo - Trondheim	64	64
Trondheim - Bod	142	136

Table 2: Total number of measurement drive runs per route per operator.

back-end and then import it into a database. Along with the measurement results, each node also provides context information (i.e. metadata) that is very valuable during the analysis. Furthermore, we access the train GPS location from the NSB system. For the coverage analysis and its implications, we use the combination of measurement data (network metadata, UDP ping and HTTP downloads) and GPS data results, which we explain in detail in the following sections.

GPS Dataset. We collect the GPS location data from the train fleet management system to identify the location of NNE measurement nodes and train's speed during the measurements. The trains update their GPS locations every 10 to 15 seconds in the NSB fleet management system.

Geo-tagged (Metadata) Dataset. The NNE nodes monitor various metadata types including the RAT, which can be *No*

³The train routes are: Oslo-Voss, Oslo-Stavanger, Oslo-Trondheim and Trondheim-Bodø.

service, 2G, 3G or 4G/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.

We record the metadata when there is a change in the value of any of these variables, for example, technology change, cell/LAC change, etc. The device pushes the updates to the measurement system when changes occur. Moreover, in order to ensure that we always have the latest information about the connection status in case the modem happens not to inform the changes in the metadata, we also pull this information from the device by querying the device periodically. More precisely, at the beginning of every minute, we record the values for all above-mentioned variables.

For our analysis, we specifically require geo-localization of the coverage information from the modems. To obtain this, we merge the metadata from the device with the GPS information we retrieve from the train. The trains update their GPS locations every 10 to 15 seconds in the NSB fleet management system. For each GPS point, we find the corresponding RAT value (e.g. No Service, 2G, 3G or 4G) from the metadata readings based on the timestamp values. We start from the timestamp of the GPS reading and define a one minute time interval prior to the GPS timestamp to look for the latest update on the RAT value. Note that if there is an update on the RAT value in this 1 minute interval, we take the latest update, and this value corresponds to the RAT value of the GPS reading. As we previously explained, besides receiving push notifications from the device reporting any RAT changes, we also pull information from the device every 1minute. Therefore, if there are no RAT updates, we used the last pulled RAT value, and this value corresponds to the latest known RAT value prior to the GPS timestamp. We define a geo-tagged data point as the GPS point where we can find a corresponding metadata reading. Note that there are cases when we do not have metadata information for the GPS reading (e.g., when the modem is down, or the IP address is lost). We discard from our dataset the GPS points with missing metadata information. These geo-tagged data points forms our geo-tagged dataset.

UDP Ping Dataset. To measure network performance and capture basic quality of service (QoS) metrics (e.g., packet loss, latency), 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 from the server within one minute.

HTTP Downloads Dataset. In order to measure application performance, we run periodic HTTP downloads using *cURL*. This choice is motivated by the fact that the majority of the Internet traffic in general and MBB in particular is HTTP traffic, which transfers over TCP port 80. Additionally, MBB operators are increasingly using transparent proxies that split and accelerate TCP connections, especially for web traffic [5]. Therefore, in this study, along with TCP port 80, we also consider a different port, namely TCP port 85, to understand the effect of these

proxies on the application performance. In order to ensure the different ports experience similar channel conditions, we perform two simultaneous HTTP downloads of size 4MB on the two different TCP ports for each of our connections once every hour.

We choose this file size to capture two relevant use cases. First, the case of business passengers who would like to use the commute as part of their work day. These users are likely to engage in downloading and sending email attachments that are several Mbytes in size. Second, the case of commuters who are accessing popular video streaming services for entertaining. DASH-based streaming, for example, segments video into small chunks and send them over HTTP. For a decent streaming quality, the size of these chunks is a few Mbytes [6]. Note that we only start a HTTP download if there is coverage. Hence, this test complements our coverage measurements by examining whether coverage actually translates into a usable connectivity that lasts for a reasonable duration to allow users completing a small task. For example, downloading a 4MB file only takes half a minute with an average speed of 1Mbit/s.

For each HTTP download, we log the start time of the download and set the timeout to 900s from the start time. Additionally, we collect information on the time to first byte (TTFB), the total download time and the transfer size (as a proportion of the 4MB target file that successfully transferred during the download time). We also log the HTTP error code and utilize it in order to deduce whether the download was successful or not. We collect this dataset from March until mid-May 2015, which overpasses the time interval we use to collect coverage measurements by 6 weeks. This allows us to analyze HTTP performance assuming that the coverage profiles are stable.

For each MBB provider and for each of the two ports, we further map the HTTP download data-point to the corresponding GPS location. To this end, we match the GPS dataset from the NSB fleet management system with the results of the HTTP measurements from the NNE mobile nodes. For each HTTP download, we define the download interval and we find the GPS point that is closest to the download start time and falls within the download interval based on the timestamps. This allows us to geo-locate the HTTP download start time.

3. Identifying Coverage Profiles

Our goal is to build a *coverage mosaic*, where we segment the routes and classify the coverage of each segment into coverage profiles that capture the distribution of RATs as the end-user would experience it. The purpose of building such a coverage mosaic is to enable further analysis in terms of network performance characterization and reliability. In this section, we propose and evaluate the use of hierarchical clustering to characterize coverage patterns in space and time. However, this comes with a number of challenges, which we formulate below.

3.1. Motivation and Problem Formulation

Investigating coverage patterns in terms of distribution of different technologies in the same area over time is challenging for



Figure 3: ECDF of train speed over the segments of each of the 4 routes for a) Telenor and b) Telia.



Figure 4: ECDF of the number of geo-tagged data points per grid block per route for a) Telenor and b) Telia.



Figure 5: CDF of the number of train runs (i.e., measurement repetitions) on each route segment with at least 4 GPS points per route. We measure different routes using hardware devices that operate depending on the schedule of the passenger trains. The Oslo-Trondheim route has the lowest number of repetitions, while the Oslo-Stavanger and the Trondheim-Bodo route has the highest number of repetitions. In order to have comparable coverage charts for the route segments we classify, we analyze only the ones where we collected data during at least 75% of the maximum of measurement repetitions. The vertical lines we show in the plots represent the 75% of the maximum number of measurement runs (on the x axis) we indicate in Table 2. The lines are color-coded to match the Train Route legend. From the intersection of the ECDF with the vertical lines, we discard from our analysis the route segments at the left of the vertical lines that have less than 75% of the maximum possible measurement repetitions per route and operator.

three reasons. First, the RAT distribution varies greatly from one segment to another based on the deployment of base stations in an area. This information is usually not available from an objective source. Additionally, connectivity upgrades are common and operators do not make their strategies public. Second, the geo-tagged dataset is difficult to work with because of its large volume and spatio-temporal inconsistency (i.e., data point's location and time of reading differ from run to run over the same route). Third, the measurement data is noisy because of a number of factors, including specific geography of the area, variable train speed, number of passengers in the train or congestion in the network. For example, all the end-users active in an area at a moment in time might not simultaneously use the fastest available RAT. All these reasons make efficient characterization of the network coverage using statistical descriptors challenging and cumbersome.

To tackle these challenges, we design, implement and evaluate a machine learning methodology that can help us characterize coverage patterns in the areas of interest. Our methodology contains two separate parts: *data morphing* and *clustering*. First, in the data morphing, we propose the segmentation of the region of interest in smaller areas and spatially group the geo-tagged data points in these segments. Second, we focus on capturing the prevalent coverage profiles and classify each route segment to a profile accordingly.

Although good or bad coverage may seem like a straightforward classification, due to the challenges we present above, it is surprisingly hard to give predetermined quantitative definitions of what is good or bad coverage when focusing on the combination of different RATs in the same area (e.g., is continuous 2G bad coverage? is it worse having intermittent 3G coverage? how can one set the thresholds on the RAT distributions to classify good or bad coverage?). We expand on this issue in the following Section 3.3. We propose the use of unsupervised clustering to help us with the classification. The spatio-temporal heterogeneity of the geo-tagged dataset we collect makes the clustering algorithm a good fit for identifying patterns in the coverage offered by MBB providers. In particular, we choose the hierarchical clustering algorithm [7, 8], which clusters data instances based on their similarity [9], thus highlighting the prevalent coverage profiles in the region.

3.2. Data Morphing

Our geo-tagged dataset from repetitive runs consists of numerous time-stamped instances of network-specific variables at different geographic coordinates along the railway routes. We identify the objects in our dataset as categorical variables (i.e., the RAT value) with dynamic location (i.e., the results from the measurement device does not always come in the same point in space). The interaction between the spatio-temporal dimensions of the dataset dictates the complexity and challenges in moving from acquiring the data to drawing knowledge through data analytics. In order to address these challenges, we begin by organizing the dataset into instances that we can easily compare.

Spatial binning. We first divide the railway routes into smaller segments using a fix grid of $2km \times 2km$ tiles that we superimpose on Norway's map. Each square grid block that overlaps on the train routes contains a segment of the route. The resulting *segments* are disjoint and uniquely identified by the fix

spatial coordinates of the square *grid blocks* that contain them. We then partition the geo-tagged dataset by grouping the data points that fall along the same route segment.

In order to make an informed decision on the size of the grid we use for spatial binning, we first investigate the speed distribution of trains, which we depict in Figure 3. We show both operators to illustrate that the speed of the train affect both geotagged datasets in the same way. We observe that majority of speeds are between 75-100 kph and in 98% of the time the train stays below the maximum speed of 120 kph for all the routes.

Recall that the granularity of the GPS data is 10-15 seconds and in order to have a long enough period to sufficiently capture the data-plane performance, we need multiple GPS readings within the same bin. Based on the speed distribution, the smallest possible grid block size that captures at least one GPS points even at high speeds is $500m \times 500m$. However, a single point per grid block is not enough to allow performance analysis. Considering all the above observations, we decide to use a grid block of size $2km \times 2km$.

In order to validate our choice, we illustrate the total number of GPS points within the grid blocks which the train traverses in Figure 4. As expected, due to variation in the speed and the variation of the route segment lengths per grid block, the number of GPS points varies. More specifically, for the Oslo-Trondheim and Trondheim-Bodø routes approximately 75% of the route segments have 4 or more geo-tagged data points for both operators. For the Oslo-Stavanger and Oslo-Voss routes the percentage of route segments with 4 or more geo-tagged datapoints drops to approximatively 65% for both operators. We leave for future work a detailed analysis of the impact that the size of the grid we use for spatial binning has on the coverage patterns we observe. The analysis of the spatial data in terms of spatial sampling patterns is an important and hard question, and we are currently working on informing spatio-temporal sampling guidelines based on geographical information system (GIS) knowledge for network measurement platforms, such as NorNet Edge or the upcoming MONROE⁴ platform. As part of our future work, we will also consider the technical guidelines of Data Specification on Geographical Grid Systems⁵ and use INSPIRE reference grid for spatial binning in order to ease the merging of our (binned) results with other potential data sources.

Coverage Chart Time Series. After the spatial binning, the route segment with fix spatial coordinates becomes the object that we further characterize in terms of mobile coverage. A route segment is characterized by a variable number of RAT readings, corresponding to the geo-tagged data points from every run that the $2km \times 2km$ area encloses. In this second step, we transform the categorical variable representing the RAT at each geo-tagged data point into a set of continuous variables that show the distribution of each of the RAT values (i.e., 4G, 3G, 2G or noS) over the set of data points along a segment of

⁴https://www.monroe-project.eu/

⁵http://inspire.ec.europa.eu/documents/Data_

Specifications/INSPIRE_DataSpecification_GG_v3.1.pdf



Figure 6: The distribution of RAT per route segment for Telenor and Telia along Oslo-Stavanger train route. The x-axis represents the route segment ID and the y-axis represents the percentage of each RAT for that particular segment using median, first and third quartiles. The points on the figure are the median/quartile values, while the lines show the fitted curves to these points.

the route. We define the distribution of RAT over one run as the segment's **coverage chart**. For example, if a segment contains 5 different geo-tagged data points [noS, 2G, 3G, 3G, 4G], then we can derive the coverage chart of the segment for the measurement run: 2G: 20%, 3G: 40%, 4G: 20%, noS: 20%.

We merge the runs independently of the train trip direction along a route to generate a coverage chart time series. Note that due to variations of train speed and the variations on where the GPS points are sampled, the number of geo-tagged data points per segment may vary. For example, we might have 3 geotagged data points in one run and 4 geo-tagged data points in the consecutive run for the same grid. As discussed in the previous section, our analysis requires a minimum number of 4 GPS points per route segment per run for each route and operator. In Figure 5, we illustrate the empirical cumulative distribution function (ECDF) of number of train runs that generate at least 4 GPS points on each route segment per route and per operator. We measure different routes using different nodes and depending on the schedule of the passenger trains that the nodes are placed, each route has a different number of runs. For example, the Oslo-Trondheim route has the lowest number of repetitions, while the Oslo-Stavanger router has the highest number of repetitions. We observe that by selecting 75% of the total runs per route, we have sufficient number of runs per each route to do the classification. The dashed vertical lines on the ECDF plots are color-coded to match the train routes legend and represent the 75% thresholds for the number of measurement repetitions we require to build the coverage chart time series of each route segment. Hereinafter, in order to have comparable coverage chart time series for the route segments we classify, we analyze only the route segments where we collected at least 4 geo-tagged data points for at least 75% of the total of measurement repetitions. We discard the route segments at the left of the intersection of each vertical line with its matching ECDF curve. The proportion of route segments we discard per router per operator varies between 37% (e.g., for Telenor on Oslo-Trondheim or Telia on Oslo-Voss) and as much as 62% (e.g., for Telia on Trondheim-Bod).

3.3. The Clustering Approach

After morphing the dataset, we reduce the problem to the matter of quantifying the similarity between the segments' coverage charts. For the Oslo-Stavanger route, we illustrate the spatial variation of the distributions of different RATs in Figure 6a for Telenor and in Figure 6b for Telia. We immediately observe that, for both operators, the distribution of RATs greatly varies in the spatial domain, supporting our claim that defining thresholds on some statistical descriptors of the RATs distributions to profile coverage is difficult and non-adaptive.

In this section, we present the similarity metric we choose, the clustering method we follow and the approach we use to determine the optimal number of coverage clusters. We perform the clustering of segments on a per-route basis, thus applying the same methodology on datasets we collect along 4 different routes.



Figure 7: Clustered Time Series on the Oslo-Stavanger route, for a) Telenor and b) Telia. Each subplot contains in the upper part - the dendrogram with the hierarchical clustering results, and in the lower part - the measurements tile plot (i.e., the mosaic containing all the measurements in the dataset) arranged following the order of elements we show in the dendrogram. The dendrogram illustrates how the grid blocks we represent on the x axis group together according to the extended Jaccard distance metric we represent on the y axis of the plot. Furthermore, it shows how homogeneous the resulting clusters are in terms of the distance between the grid blocks within. The tile plot illustrates all the measurements we collect on the Oslo-Stavanger route and has on the y axis the time component (i.e., the drive runs the train performs on the route in chronological order from down to top), and on the x axis the spatial component (i.e., the grid blocks we use to segment the Oslo-Stavanger route). Each tile of the mosaic plot represents one measurement instance of the RAT tuple [2G, 3G, 4G, noS], color-coded per level (e.g., green for 4G, orange for 3G, blue for 2G and red for noS) and with a color gradient proportional to the value of the distribution of the connection mode.



(a) Telenor (b) Telia Figure 8: Characterization of the coverage profiles we derive using all the data we collect during the 5-months measurements period along the four measured routes in Norway, for (a) Telenor and (b) Telia. The coverage profile consists of four features we label on the x axis (%2G, %3G, %4G, %noS) that show the distribution of the four RATs over the segments in the same cluster and whose values we illustrate on the y axis.

Similarity metric. In order to calculate the similarity between two segments, we organize the coverage time series into vectors of coverage charts we measure at each run over a route. The length of the vector is fix and equal to the number of measurement runs we register on a route. In the case where the coverage chart for a segment for a run is missing (either due to hardware issues during the run or due to sifting the data based on the minimum required number of data points), we populate the coverage chart with null value or the corresponding coverage mode variables. We then use the extended Jaccard measure [10] to evaluate the similarity between two objects. In Section 4.2, we investigate the similarity in the measurements from repeated runs and quantify the number of minimum measurement samples which we can use to characterize each route segment in terms of coverage as experienced by the end-user. In Section 4.4, we then generate the per-route-segment coverage chart using a fix number of non-null samples. This further allows us to understand the impact of populating the time series of a route segment with null values for the runs when we were unable to generate the coverage chart.

Clustering method. The clustering method we select is an average-link hierarchical clustering method, which organizes the data objects into a multi-level structure, based on the similarity between objects. Such methods consider the distance between two clusters to be equal to the average distance from any member of one cluster to any member of another cluster. More specifically, we employ here Ward's minimum variance method for clustering [11], which aims at finding compact, spherical clusters.

Optimal Number of Clusters. Hierarchical clustering leaves the task of detecting the optimal number of clusters to the user. To this end, we evaluate a *validity index* for different number of clusters [12]. The number of clusters that generates the best value for the index is then chosen as optimal. There is no general consensus on which validity index should be used. In this paper, we consider the Silhouette index [13] that represents an average, over all the clusters, of how similar the data in each cluster is.

4. Clustering Results

In this section, we run the proposed methodology to characterize coverage along the four Norwegian train routes and then we analyze the results.

4.1. Coverage Clusters Analysis

We apply the hierarchical clustering method separately to the coverage chart time-series of the segments covering each of the train routes. In other words, for each route, we calculate the similarity between segments' coverage chart time series and we group together the segments with similar coverage patterns. We then determine the optimal number of coverage clusters using the Silhouette index. For all the four routes and both operators, the Silhouette index gives 2 clusters, which we detail next.

Prevalent coverage profiles. In Figure 7, we depict the results for the hierarchical clustering obtained, both for Telenor and Telia. Each subplot contains in the upper part the dendrogram of the clustering results, and the tile plot of the coverage clusters of grid block time series in the lower part. The optimal clustering of route segments defines the dominant *coverage profiles* for the two MBB operators. We observe that the clustering algorithm identifies two main coverage profiles, which we generically label as coverage *Profile A* and *Profile B*. In Figure 7, the color-coded values of the tiles show that route segments with similar coverage chart evolution are grouped together.

In Figure 8, we illustrate the characteristics of these two coverage profiles, namely the average RAT distributions [2G, 3G, 4G, noS] over all runs for the segments in the same cluster. When analyzing the coverage profiles, we note that in the areas with *Profile A*, Telenor has around 70% 3G accompanied with 15% 4G, while Telia compensates with higher 3G availability (85% 3G) for its slower 4G deployment (5% 4G). This clearly shows the different deployment strategies of the operators. Furthermore, we observe slight differences in the profiles of different routes. For example, the Trondheim-Bodø route in the *Profile A* areas of Telia clearly stands out because it has a lower 3G and higher noS percentage compared to the other routes. For the *Profile B* coverage areas, we observe a high degree of *No Service* for both operators, which combines with a smaller ratio of 3G and 2G. The distribution of 2G and 3G varies considerably among different routes.

Coverage profiles on route. In Figure 9, we show the spatial distribution of the two coverage profiles along the four train routes. We observe that Telenor had a slightly larger area with *Profile A* coverage than Telia. For both operators, we observe that coverage *Profile B* is extensive along the critical transport infrastructure. This is mostly due to the very challenging geography of Norway where there are many mountains that the trains need to traverse. Furthermore, in order to minimize the difficult mountain crossings during winter, many tunnels have been built along the railroads. Moreover, in contrast to other European countries, the rural areas in Norway is very thinly populated (see Figure 2). Therefore, part of the railroad has known to have no coverage or only 2G coverage and these are the areas where an infrastructure improvement is needed.

Next, we analyze the coverage profiles in more detail for the Oslo-Stavanger route. In Figure 10, we depict the results for the hierarchical clustering we obtain for Oslo-Stavanger, both for Telenor and Telia. Each subplot contains the dendrogram of the hierarchical clustering grouping according to the similarity measures we choose. We group the segments into two main coverage clusters, namely *Profile A* and *Profile B* coverage clusters. We note that for both operators, the two clusters are well distanced. Additionally, the distances between elements within the same cluster are relatively small. Thus, the dendrogram allow us to clearly observe the separation between the clusters according to the similarity measure, validating the result of the validity index.

Though, using the Silhouette index, we systematically discover two prevalent coverage clusters that dictate the coverage profiles, we can further divide these clusters into smaller sub-clusters with more homogeneous profiles. In some cases, other indexes such as the Dunn or DB indexes indicate a larger number of clusters, because of the heterogeneity between the coverage time series of each of two major clusters. We see in Figure 10 that the coverage clusters contain several coverage sub-profiles that highlight the predominance of one RAT or the mixture of several RATs. For example, in the case of Telenor, we identify 4 different sub-clusters in the *Profile A* "good" coverage cluster, underlining the increasing heterogeneity of technologies in the MBB networks. For future work, we plan to collect GPS with a 1 second granularity and further analyze the impact of sampling on the clustering results.

4.2. Coverage Clustering Stability

In this section, we focus on coverage cluster stability and investigate the minimum number of runs (i.e., different samples in time) that is sufficient to classify a segment in one of the



Figure 9: Map of the Coverage Mosaic we derive by clustering and analysing the route segments using all the measurement repetitions we collect throughout the 5 months measurement interval. The route segments are characterized by the two different coverage profiles (their color and shape shows the coverage profile to which each was assigned), for (a) Telenor and (b) Telia.



Figure 10: Clustering dendrograms showing the prevalent coverage profiles (i.e., Profile A and Profile B) and the coverage sub-profiles along Oslo-Stavanger route, for (a) Telenor and (b) Telia. We generate the dendrogram of the hierarchical clustering according to the similarity measures we show on the y axis. On the x axis we show the ID for each route segment we cluster.

main coverage profiles, namely "Profile A" with good coverage and "Profile B" with bad coverage. Our goal is to quantify how much additional information regarding the coverage can each drive run bring to the clustering problem and when the classification of a route segment with one of the above-mentioned profiles is stable.

To this end, we run the clustering approach we explain in Section 3 for varying number of runs n, where n is between 1 and 50. Using the resulting coverage chart time series from nmeasurement runs, we cluster the segments and separate them in Profile A and Profile B coverage clusters, which we have previously established to intuitively stand for "good" and, respectively, "bad" coverage. We mention that we use here as input the results of the prior analysis in Section 4.1 in terms of the existence of the two main coverage clusters, but the clusters resulting after each iteration are obviously different from the ones we obtain when using all the available dataset. We use the profile characteristics we illustrate in Figure 8 to guide the labelling of the coverage profiles at each iteration into "Profile A" or "Profile B", so we are always trying to understand whether a segment had "good" or "bad" coverage. For example, in order to find the clusters of segments using only 2 measurement runs (n = 2), we select any pair of runs among all runs and apply the clustering algorithm. Thus, we obtain an assignation of coverage profile per route segment, for all routes and both operators. For each n, we repeat this exercise for 100 different combinations of n runs out of the ones available on each route. We obtain 100 different assignation of the coverage profiles per segment for each size n of the set of runs we use as input. In order to gauge the differences between the 100 coverage profile assignments for every n input measurement runs, we calculate the similarity between these 100 coverage profile assignments. We use the Jaccard distance [10], which is well-defined for binary vectors (for each segment, we assign 1 for *Profile A* and 0 for *Profile B*).

In Figure 11 we show the average Jaccard distance between the coverage profile assignments as a function of the number of runs we use as input. We show this for each operator and every route we measure. We conclude that individual runs are highly dissimilar, generating highly variable assignations of coverage profiles for the analyzed grid blocks. This is reflected in Figure 11 with the value of the average distance corresponding to the number of runs equal to 1. However, we observe that, in order to obtain a stable coverage profile assignment to each route segment, the minimum number of drive runs required is between 5 and 10. This result is consistent over all the routes, for both operators.

4.3. Coverage Profile Adaptability

Previously, we determined that 5-10 measurement runs bring enough to decide whether a route segment belongs to Profile A or Profile B. In this section, we aim to capture the evolution over the measurement period for these two coverage profiles in terms of the average distribution of RATs over the route segments that fall within each coverage cluster. Here, we assume we have a set of consecutive runs per segment. We then use a sliding window of 10 measurement runs and employ the clustering approach we propose in Section 3 to assign each segment to the good or bad coverage cluster. In other words, clustering is done based on the data collected over 10 consecutive runs per segment where the first run is shifted in time. For example, the first clustering results considers the first 10 runs in a segment while the second clustering results consider 10 runs starting from the second run in the same segment. This analysis allows us to capture the technology upgrades over the period of 5 months we measure. In Figure 12, we exemplify this analysis for the Trondheim-Bodø route, which is the one where we collect the highest number of measurement runs. We note that the coverage profiles are overall stable for both operators. However, for Telenor we observe a slight increase in the 4G distribution in the areas with Profile A coverage. Also, there is a small improvement in the 3G distribution in the areas with Profile B coverage along the Trondheim-Bodø. More specifically, in the first 20 repetitions of deriving the RAT distribution in the areas with Profile A, the average 4G distribution is 14.8% and in the final 20 repetitions, the average 4G distribution increases to 18.6%. Similarly, in the areas with Profile B coverage, we note that the initial average 3G distribution of 8% increases slightly to 9.8% over the last 20 repetitions of coverage profiling using a temporal sliding window of 10 runs. This shows that our methodology is capable of capturing technology upgrades in the operators' networks.

4.4. Coverage Clustering Optimization

In the previous section we show that the similarity between different measurement runs is high and to obtain a stable coverage profile assignment to each route segment, we can reduce the number of distinct drive runs to as little as 5 to 10 runs. In this section, we re-run the clustering analysis of the route segments for each route using only the latest 10 different measurement runs per route segment. We then analyze the coverage profiles and investigate whether the coverage profile assignment changes significantly compared to when we were using all available measurements.

For each operator, for each route and for each route segment we select where applicable the results from the latest 10 measurement runs. Using this new dataset, we re-run the same clustering algorithm, using the extended Jacccard distance as similarity metric, the Ward grouping algorithm to cluster route segments and the Silhouette index to derive the number of coverage profiles. We note that the measurement runs that provide the coverage samples may differ for distinct grid blocks. As we previously note in Section 3.3, during some measurement runs, the coverage chart for a route segment might be missing because we were unable to characterize the coverage along the route segment. This may happen either because of insufficient measurements along the route segment to generate a coverage chart (according to the thresholds we impose), lack of coverage and subsequent impossibility of measuring with our mobile devices or device issues. Considering the latest 10 samples of the coverage chart for a route segment irrespective of the run allows us to avoid the artificial population of the coverage chart with null values. We continue calculating the similarity between route segments using the extended Jaccard distance. Based on this, we then group the elements and observe the coverage clusters that further dictate the prevalent coverage profiles.

We illustrate in Figure 14 the characteristics of the coverage profiles that correspond to the clusters of route segments. When comparing these results with the coverage profiles built using all the available data, we observe similar coverage profiles. Especially on the Oslo-Voss, Oslo-Stavanger and Trondheim-Bodø routes for both operators the coverage profiles are remarkably similar. We draw the same conclusion for Telia on the Oslo-Trondheim route. However, in the case of Telenor's MBB service along the Oslo-Trondheim route, we note that the route segments with Profile B coverage have comparable noS/2G/3G/4G distributions within the profile, indicating that along these segments we recorded frequent RAT changes from the measurement device. Unlike for the route segment along the same route with Profile A coverage, where 3G is obviously dominant with 80% ratio, in this coverage profile there is no clear dominant RAT along the route segments, even if noS has the highest ratio of more than 35% in average. Despite the fact that we register quite significant 4G presence in Profile B, the coverage is patchy.

This difference from the previous clustering result comes from the fact that by considering the route segments along the route with exactly 10 different measurement repetitions, we are able to analyze a larger proportion of the route but discard the end-user experience by not accounting for the cases when the device was not able to measure any RAT. In Figure 13, we illustrate on a Norway map the route segments we were able to analyze and profile using a coverage chart time series derived from the latest 10 measurement runs on each segment indepen-



Figure 11: Stability of coverage profile assignment in function of the number of measurement runs we use to build the coverage chart time series of the route segments.



Figure 12: Evolution of the coverage profiles when using a sliding window of 10 measurement runs to derive them.

dently. When comparing with the coverage mosaic from Figure 9, we confirm that, indeed, the total number of route segments we are able to analyze significantly increases. By discarding the measurement instances in some drive runs where we were unable to generate a coverage chart for a route segment, we are essentially discarding the cases when the end-user is not able to connect to the network. Though for the other routes this does not impact the coverage profiling, in the case of the Oslo-Trondheim route the coverage patterns seem to shift, as there is small amount of no service that we measure along the route (compared to the other routes). Thus, along this route, the patterns in the coverage evolve from dominant-3G (Profile A) and dominant-noS (Profile B) to dominant-3G (Profile A) and patchy-coverage (Profile B) with frequent RAT changes.

For the route segments we previously characterized using all the coverage measurements collected throughout the 5-month measurement interval, we investigate whether the coverage profile assignment changes (i.e., the coverage) when we considered only 10 different measurement samples along the routes. In other words, we investigate to which degree the route segments from the coverage mosaic we depict in Figure 9 change coverage profile assignment in the coverage mosaic we depict in Figure 13. In Table 3, we summarize the number of route segments that change the coverage profile after reducing the number of measurement runs to only 10. The fraction of route segments that change coverage profile assignment varies between 5% (routes like Oslo-Voss or Trondheim-Bod), to as much as 17% (routes like Oslo-Trondheim for Telenor). As noted above, this is due to the change in the definition of coverage profiles along the route after discarding the cases where we were unable to generate the coverage profile of some route segments.

Route	# Route Segments (Proportion	
	Telenor	Telia
Oslo - Voss	4 (5.6%)	9 (10%)
Oslo - Stavanger	12 (7.8%)	13 (7.14%)
Oslo - Trondheim	33 (17%)	24 (12%)
Trondheim - Bod	12 (5%)	19 (12%)

Table 3: The number of route segments that change Profile Coverage assignment when using the latest 10 measurement runs per route segment (without synchronization between runs over the route segments). For each value, we show in parenthesis the proportion of the total number of route segments these ones represent.

5. Coverage Implications: Reliability and Performance

In this section, we focus on how coverage profiles correlate with the performance of the networks from the end-user point of view. During the performance analysis, we use the profiles we derive using all the data (see Section 4.1) and we investigate the performance throughout the whole measurement period. This analysis opens the future possibility of using the resulting coverage mosaic with coverage profiles as an indicator for network performance. Furthermore, this will enable the



Figure 13: Map of the Coverage Mosaic we derive by clustering and analysing the route segments using only the latest 10 measurement repetitions we collect in the 5 months measurement interval. The route segments are characterized by the two different coverage profiles (their color and shape shows the coverage profile to which each was assigned), for (a) Telenor and (b) Telia.



Figure 14: Characterization of the coverage profiles we derive using the latest 10 measurement runs per route segment along the four measured routes in Norway, for (a) Telenor and (b) Telia. The coverage profile consists of four features we label on the x axis (%2G, %3G, %4G, %noS) that show the distribution of the four RATs over the segments in the same cluster and whose values we illustrate on the y axis.

design of contex-aware algorithms to improve the application quality of experience for end-users. We now turn to investigating per-profile MBB performance in terms of downtime and packet loss as key performance measures.

5.1. Uptime

To calculate connection uptime in a time window T, we divide the number of sent packets in T by the length of T. In

our case, T represents the time that the NNE node spends inside a grid block, and has a minimum value of 30-45 sec. Figures 15a and 15b show the CDF of the fraction of uptime for each route and cluster combination. As expected there are clear differences between areas with different coverage profiles. For example, areas with *Profile B* coverage ("bad" coverage) exhibits very low uptime caused by the high percentage of no service. Furthermore, there are clear spatial differences be-



Figure 15: Uptime per route, run and cluster. There is a tight coupling between uptime and the coverage profiles we show in Fig. 8.



Figure 16: The distribution of packet loss in a grid block for different routes.

tween operators. For instance, the worst performing route in the *Profile B* coverage cluster is Oslo-Stavanger for Telenor and Trondheim-Bødo for Telia. Trondheim-Bødo is the best performing route with *Profile B* coverage for Telenor. These differences indicate that multi-connectivity, i.e. the use of several operators simultaneously, can improve users experience along the same route. Analyzing differences in uptime between operators and between routes in conjunction with routes coverage profiles (see Fig. 8), indicates a tight coupling between coverage profiles and uptime.

5.2. Packet loss

Figure 16 shows the CDF of packet loss in a grid block. We measure a significant difference in the extent of loss between *Profile A* and *Profile B* coverage clusters. Loss, however, remains high in the areas with *Profile A* ("good" coverage) coverage. Between 30% and 50% of our samples, depending on the operator and route, exhibit more than 1% packet loss. This can be explained by the co-existence of different RATs and the need for frequent handovers in the areas with *Profile A* coverage. Further, we observe that the ranking of route segments with *Profile A* coverage in terms of packet loss matches their ranking in terms of uptime. There is, however, less similarity between the *Profile B* coverage routes loss and uptime rankings:

14

the worst route is always similar. We believe this similarity in ranking is because packet loss is usually experienced in areas with challenging coverage conditions (i.e., larger percentage of no service) which can also lead to a connectivity loss. We also measure how loss in a grid block varies in different runs and find that irrespective of the operator and route, the standard deviation of packet loss is twice as much the mean for at least 50% of the grid blocks. Interestingly, this variability is higher for areas with "Profile A" coverage (which can be perceived as "good" coverage), which underlines the fact that one-off measurements are not enough to make conclusions about performance under mobility.

6. Coverage Implications: Application Performance

In this section, we continue the analysis from Section 5 on the implications of the coverage profiles and investigate how the latter correlate with the application performance as experienced by the end-user. In particular, we focus on HTTP performance and we study TCP port 85 along with TCP port 80 in order to understand the effect of the proxies on the application performance. We employ the HTTP Download dataset we describe in Section 2.2. In order to identify the corresponding route segment for an HTTP download data-point, we use the GPS information of the HTTP download data-point to identify the unique grid-block that contains the GPS location of the HTTP data-point. Thus, by identifying the grid block, we also retrieve the corresponding train route segment delimited by that grid block and its coverage profile.

Next, we first describe how we identify the presence of a proxy in one of the operator's networks. We then analyze the HTTP performance in terms of download success and failure. Finally, we characterize the HTTP downloads based on the features we collect for each data-point.

6.1. Detecting the presence of a proxy

Assessing TCP and HTTP performance is not a trivial task since there can be several middleboxes on the end-to-end path that modulate TCP behavior. Transparent web proxies can alter the end-to-end communication in several ways that include content caching, traffic redirection, object rewriting (e.g. image compression), and connection lifecycle manipulation. Since we control both end points (i.e. the HTTP server and client), we can detect the presence of transparent proxies for web traffic by inspecting packet traces of a complete HTTP transaction. Note that proxies' behaviors like traffic redirection and object rewriting are not relevant to the test scenario because we are exchanging immutable content from a single server. Further, we find that our HTTP server receives an HTTP request and engages in a packet transfer whenever a download is attempted meaning that neither operator caches the requested file. To test for presence of connection lifecycle manipulation proxies, we perform the following:

- Check if the initial TCP handshake is delayed: Transparent proxies are known to delay the initial handshake until the client sends an HTTP GET request [5].
- Inspect TCP roundtrip time measured at the server side: Presence of a connection splitting proxy pretending to be the other end of the connection would result in TCP roundtrip times that are not consistent with the path end to end delay.
- Cross-match all TCP messages received by the server side with all messages sent by the client: A mismatch in these messages confirms the presence of a transparent proxy.

We perform these preliminary tests both for Telia and Telenor. We find evidence that Telia uses a transparent proxy. For Telenor, we cannot confirm nor preclude the use of a web proxy. More specifically, we find that Telia's proxy delays the initial TCP handshake until the client sends an HTTP GET request. We also find that server side TCP roundtrip times are mostly less than 5 msec, which is far lower than a typical RTT in a MBB network. Finally, we find a clear mismatch between TCP messages received by the server side with all messages sent by the client. The server received a receiver window full messages that were not sent by the client. By doing this, we believe that the proxy was attempting to slow down the server without reducing the size of its congestion window. Note that this evidence is only found when accessing content over traditional web ports (e.g. ports 80 and 8080). We detail in the following section the impact of the web proxy in both coverage profiles.

6.2. HTTP Download Success/Failure

In Figure 17, we show the breakdown of the HTTP dataset based on the coverage profile of the route segment where each download initiated. We observe that there is a large difference between the sample sizes of HTTP downloads in Profile A and Profile B. For Telenor, there are 8 times less data-points along Profile B route segments than in Profile A route segments. Similarly, for Telia there are 5 times less data-points initiated in Profile B route segments. This is reasonable though since, based on the coverage profiles' characterization in Section 4, Profile B has a large degree of *No Service* in the RAT distribution. Thus, it is likely that several of the HTTP downloads we schedule in the areas with Profile B coverage cannot start because of the lack of connectivity.

Moving downwards through the HTTP tree in Figure 17, we observe the separation for each coverage profile, based on the operator (Telenor/Telia) and the port number (port 80/85). For each sub-branch, we then break apart the successful/failed downloads and calculate the corresponding number of route segments (RS). A download is labeled as failed, if it does not conclude with transferring the whole 4MB file, or if cURL exits with a non-zero code. The later happens when cURL fails to complete the initial handshake, or when network connectivity is lost while downloading the file. We observe that in the case of HTTP downloads that start along a route segment with Profile A coverage, the rate of success is very high. While we observe comparable high success rate in both ports (84%) for Telenor, we note that for Telia the rate of success in port 80 (92%) is significantly higher than in port 85 (85%). Given that the measurement setup is such to ensure the same conditions for downloads in port 80 and 85. This difference in performance indicates that web content served over port 80 in Telia enjoys a differential treatment because of the use of transparent web proxies.

For the HTTP downloads that initiate along route segments that have Profile B coverage we observe, as expected, that the rate of failure is very high for both operators. Similar to the case of HTTP Downloads that initiate along route segments with Profile A coverage, in the case of Telia there is an obvious difference between the failure rate in port 80 (64%) and the one in port 85 (71%). However, the overall failure rate in Telenor (97%) is much larger than the one in Telia (70%), though the sample sizes in both cases are comparable. Apart from accounting for the benefits of using a middlebox for performance enhancements, we can get more intuition into what may be the cause for this difference by analyzing the characterization of the coverage profiles we depict in Figure 8. We observe that the coverage Profile B in Telia has a slightly higher rate of 3G and 2G than in Telenor. Also, overall we observe that the number of route segments we classify with Profile B in Telia is higher than the one in Telenor. This is apparent also from the tile plot we depict in Figure 7 for the Oslo-Stavanger route, where the Profile B cluster is 30% larger in Telia than in Telenor. Also,



Figure 17: Breakdown of the HTTP downloads dataset per coverage profile, operator (Telenor and, respectively, Telia) and per port (for each operator, we measure port 80 and port 85 at the same time).

we observe that the two clusters of route segments we use to define coverage profiles are more homogeneous and better distanced in the case of Telenor than in the case of Telia. This explains why for Telenor most HTTP downloads that initiate along route segments with *Profile B* coverage end in a failure, while for Telia as much as 30% of them succeed.

6.3. HTTP Download Characterization

As we previously described in Section 2.2, we characterize each HTTP download data-point beyond the binary description of success/failure by measuring the time to first byte (TTFB) and the transfer size (the proportion of the target file that downloaded using cURL). In Figure 18, we show the ECDF for all the three above-mentioned variables for (a) Telenor and (b) Telia. In each subplot, we break apart the ECDFs as a function of the port we measure and also the coverage profile.

In all the ECDFs for the HTTP transfer size and TTFB we observe the separation between the downloads that initiate in Profile A route segments and the ones that initiate in Profile B route segments. The variable that makes the clearest separation between the two coverage profiles for both operators is the transfer size (i.e., proportion of target file downloaded). This is due to the fact that this parameter captures until which point the HTTP download continued before losing connectivity, which is important especially in areas with Profile B coverage. For Telenor, we observe that in more than 90% of the Profile B downloads the transfer is incomplete, while for approximately 90% of the Profile A downloads the transfer is complete. This also validate the previous conclusions we drew from analyzing the HTTP downloads tree. For Telia, we note that in the case of HTTP downloads that initiate in Profile B route segments, the transfer completes for approximately 25% of the cases. This verifies our observation from the previous section that a slightly larger degree of 2G or 3G in the RAT distribution in Profile B for Telia explains the higher proportion of completed HTTP downloads. We observe only slight differences when analyzing the transfer size in port 80 and port 85, with a higher proportion of HTTP downloads that complete in port 80 than in port 85.

7. Related Work

Building accurate and reliable coverage maps has been in the attention of the community and a magnitude of work exists in this area. Drive tests are widely used by MBB operators for coverage assessment and performance monitoring. In this paper, we argue that piggy-backing mobile broadband measurements onto public transport infrastructure is an efficient, cost-effective and automated way to perform drive tests. Aside from the very high costs of drive tests, the data collected from them usually has a series of shortcomings, including variable spatio-temporal sampling and limitation of test repeatability. The drawbacks of drive tests act as incentive for the design of new methodologies that address these issues [14, 15]. In this sense, our experimental setup brings the benefit of repeatability at a low additional cost. Other apporaches including, for example, crowdsourcing platforms, may help verify coverage maps[16], but they also bring additional limitations including, for example, the lack of control on the measurement device and lack of repeatability.

Data analytics approaches are receiving much attention from the community, due to their capabilities to draw useful information from large databases collected from the network [17, 18]. Coverage prediction methodologies based on geostatistics[19, 20] in wireless networks constitute another approach in the direction of data analytics. To the best of our knowledge, this paper is the first attempt in mobile coverage profiling using hierarchical clustering of multivariate time series. Similar solutions have been proposed in the area of spatio-temporal data mining with different applications in real life e.g. [21, 22]. This technique enables us to generate adaptive coverage profiles, which are based on real measurements and reflect the deployment reality of MBB connectivity solutions and their evolution in time.

In the past years we have seen increased interest in the networking community from different parties (e.g., researchers, operators, regulators, policy makers) in measuring the performance of mobile broadband networks. There are mainly three approaches for measuring the performance and reliability of MBB networks: (i) crowd-sourced results from a large num-



Figure 18: CDF of the transfer size and time to first byte (TTFB) for all HTTP downloads on port 80 and port 85 for (a) Telenor and (b) Telia.

ber of MBB users [23, 24, 25, 26], (ii) measurements based on network-side data such as [27, 28, 29, 30] or earlier work including [31, 32] and (iii) measurements collected using dedicated infrastructure [33, 34, 35]. Network-side and active tests can be combined in the so-called "hybrid measurements" approach, as implemented e.g. in [36]. In this paper, we collect data from a dedicated infrastructure in order to have full control over the measurement nodes, allowing us to systematically collect a rich and high quality dataset over a long period of time. Unlike previous efforts that ran performance measurements, we focus on coverage and its implications in terms of network (e.g. packet loss) and application performance as experienced by end users.

Several studies focused on the causes of packet loss in MBB networks. Different groups blamed RRC state transitions [37, 38, 39, 40, 41] and showed that state demotions result in significant loss. 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 [42]. Xu et al. discussed the effect of bursty packet arrivals and drop-tail policies employed by the operators [5]. RNC-level performance analysis of UMTS networks identified correlations between RTTs and loss and their dependency on diurnal patterns and overloaded NodeBs [41]. 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 [35]. In a recent work [37], the authors conducted a large-scale measurement study of packet loss in MBB

(RRC) state transitions and the causes of a significant part of the remaining loss lie beyond the radio access network. Packet loss has also been investigated for mobility scenarios. [43] 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 [44], showing that most HTTP sessions with inter-RAT handovers are abandoned. [45] 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, our paper is the first study that ties the coverage with network reliability analysis by showing how coverage profiles can be used as an indicator for mobile broadband reliability. In this paper, we also study the implication of the coverage

networks. The study showed that a significant fraction of loss

occurs during pathological and normal Radio Resource Control

mosaic on the application performance by measuring and analyzing HTTP downloads. The analysis of HTTP allows us to investigate the presence of transparent web proxies that operators might be deploying in their networks. Performance enhancing middleboxes are widely deployed in Internet and it is of great interest to measure and characterize the behavior of them especially in MBB networks where the resources are scarce. One of the early studies in this domain investigated the web performance of different mobile Internet access technologies (GPRS, EDGE, UMTS and HSDPA) with and without the web-proxy [46]. The impact of middleboxes on measurements was explored in [47] where the authors proposed a methodology for measurements in MBB networks. Farkas et al. [48] used numerical simulations to quantify the performance improvements of proxies in LTE networks. The most thorough analysis to characterize the behavior and performance impact of deployed proxies on MBB networks was carried out in [49] where the authors enumerate the detailed TCP-level behavior of MBB proxies for various network conditions and Web workloads. Although the common belief is that the proxies provide performance benefits, Hui et al. [50] showed that proxies can actually hurt performance by revealing that direct server-client connections have lower retransmission rates and higher throughput. Wang et al. [51] showed how MBB middlebox settings can impact mobile device energy usage and how middleboxes can be used to attack or deny service to mobile devices. While these studies focus on the performance of proxies on MBB networks, they have not consider the effect of proxies on the reliability and packet loss.

8. Conclusions and Future Work

MBB networks are the key infrastructure for people to stay connected, especially in high mobility scenarios (e.g., when using public transport). MBB coverage profiling from the enduser experience while on critical public transport routes are of great importance to many stakeholders. At the same time, this is a challenging problem, since even a straight-forward classification of coverage into "good" or "bad" is very difficult to grasp in quantitative thresholds. In this paper, we evaluate the use of hierarchical clustering to build a coverage mosaic of MBB technologies in an area and analyze its implications in terms of network performance and application performance. By piggybacking network measurements onto public transportation vehicles via the NNE platform, we first obtained a unique dataset that (i) captures the coverage and performance from user's perspective and (ii) provides repetitive measurement runs on the same route, in similar conditions. Moreover, an important perk of such measurement platforms is allowing other parties, including public transport companies, to assess and compare the MBB coverage along their infrastructure to verify their service level agreement. We then leveraged hierarchical clustering in order to identify and characterize prevalent coverage profiles. Though in this study we look at the case of railways in Norway, the methodology can easily be generalized for running a similar study in other regions or applying it to a different datasets, (e.g. crowd-sourced data). A copy of the dataset we used in this paper is available for open access in Zenodo⁶, as well as the code for the clustering approach.

Our results reveal that the clustering approach can accurately group together regions with high similarity in terms of coverage. Based on the mixture of RATs and the time-domain evolution, two main coverage profiles emerge: *Profile A* -where

3G dominates, and Profile B - where No Service dominates. This maps onto the general intuition of "good" and, respectively, "bad" coverage. We then analyze the identified coverage profiles, both in terms of stability and performance. The stability analysis investigates the similarity between different runs over the same route, with the express purpose of indicating the amount of measurement repetitions we require to accurately observe stable coverage profiles. We find that we need at least 5 to 10 measurement runs in order to achieve a stable coverage profile in an area. We then focus on how coverage profiles correlate with MBB and also application performance from the end-user point of view. For this, we first assess packet loss performance per coverage profile and find that it highly varies for areas with *Profile A* coverage. This result is counter-intuitive because Profile A presents a high percentage of superior RATs. This indicates that, although we can derive this profile with few measurement runs, further characterization of the performance requires more analysis, e.g., correlation with the network congestion and measurement time of the day.

We take this analysis further and investigate the implication of the coverage profiles on the application performance, with a focus on HTTP traffic. We observe that in the route segments with Profile B coverage, the rate of failure for HTTP downloads is very high, while in the route segments with Profile A coverage, the HTTP downloads succeed with a high rate. For Telia, however, we note a rate of 30% of successful downloads even in Profile B coverage, while for Telenor this rate is very small, less than 5%. This is an artifact of the fact that the two clusters of route segments we use to define coverage profiles are more homogeneous and better distanced in the case of Telenor than in the case of Telia. While assessing TCP performance we also try to detect the presence of middleboxes that operators might deploy in their networks, such as transparent web proxies. While analyzing the characteristics of the HTTP downloads, we discover the impact of what seems to be a web proxy in Telia.

Acknowledgements

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

References

- O. B. Karimi, J. Liu, C. Wang, Seamless wireless connectivity for multimedia services in high speed trains, Selected Areas in Communications, IEEE Journal on 30 (4) (2012) 729–739.
- [2] M.-h. Han, K.-S. Han, D.-J. Lee, Fast ip handover performance improvements using performance enhancing proxys between satellite networks and wireless lan networks for high-speed trains, in: Vehicular Technology Conference, 2008. VTC Spring 2008. IEEE, IEEE, 2008, pp. 2341–2344.
- [3] A. Lutu, Y. R. Siwakoti, Ö. Alay, D. Baltrunas, A. Elmokashfi, Profiling mobile broadband coverage, 2016.

⁶http://dx.doi.org/10.5281/zenodo.47707

- [4] A. Kvalbein, D. Baltrūnas, J. Xiang, K. R. Evensen, A. Elmokashfi, S. Ferlin-Oliveira, The Nornet Edge platform for mobile broadband measurements, Elsevier Computer Networks special issue on Future Internet Testbeds.
- [5] Y. Xu, Z. Wang, W. Leong, B. Leong, An End-to-End Measurement Study of Modern Cellular Data Networks, in: Proc. of PAM, 2014.
- [6] J. Martin, Y. Fu, N. Wourms, T. Shaw, Characterizing netflix bandwidth consumption, in: Consumer Communications and Networking Conference (CCNC), 2013 IEEE, IEEE, 2013, pp. 230–235.
- [7] A. J. Scott, M. J. Symons, Clustering methods based on likelihood ratio criteria, Biometrics 27 (2) (1971) pp. 387–397.
- URL http://www.jstor.org/stable/2529003
- [8] L. Kaufman, P. J. Rousseeuw, Finding groups in data: an introduction to cluster analysis, Vol. 344, John Wiley & Sons, 2009.
- [9] H. Alt, M. Godau, Computing the fréchet distance between two polygonal curves, International Journal of Computational Geometry & Applications 5 (01n02) (1995) 75–91.
- [10] P. Jaccard, The distribution of the flora in the alpine zone, New phytologist 11 (2) (1912) 37–50.
- [11] J. H. Ward Jr, Hierarchical grouping to optimize an objective function, Journal of the American statistical association 58 (301) (1963) 236–244.
- [12] Y. Liu, Z. Li, H. Xiong, X. Gao, J. Wu, Understanding of internal clustering validation measures, in: IEEE International Conference on Data Mining (ICDM), IEEE, 2010, pp. 911–916.
- [13] P. J. Rousseeuw, Silhouettes: a graphical aid to the interpretation and validation of cluster analysis, Journal of computational and applied mathematics 20 (1987) 53–65.
- [14] Tektronix, Reduce Drive Test Costs and Increase Effectiveness of 3G Network Optimization, Tech. rep., Tektronix Communications (2009).
- [15] W. A. Hapsari, A. Umesh, M. Iwamura, M. Tomala, B. Gyula, B. Sebire, Minimization of drive tests solution in 3gpp, Communications Magazine, IEEE 50 (6) (2012) 28–36.
- [16] M. Molinari, M.-R. Fida, M. K. Marina, A. Pescape, Spatial interpolation based cellular coverage prediction with crowdsourced measurements, in: Proceedings of the 2015 ACM SIGCOMM Workshop on Crowdsourcing and Crowdsharing of Big (Internet) Data, ACM, 2015, pp. 33–38.
- [17] I. Leontiadis, A. Lima, H. Kwak, R. Stanojevic, D. Wetherall, K. Papagiannaki, From cells to streets: Estimating mobile paths with cellular-side data, in: Proceedings of the 10th ACM International on Conference on emerging Networking Experiments and Technologies, ACM, 2014, pp. 121–132.
- [18] F. Murtagh, A. E. Raftery, Fitting straight lines to point patterns, Pattern recognition 17 (5) (1984) 479–483.
- [19] B. Sayrac, A. Galindo-Serrano, S. B. Jemaa, J. Riihijärvi, P. Mähönen, Bayesian spatial interpolation as an emerging cognitive radio application for coverage analysis in cellular networks, Transactions on Emerging Telecommunications Technologies 24 (7-8) (2013) 636–648.
- [20] H. Braham, S. Ben Jemaa, B. Sayrac, G. Fort, E. Moulines, Low complexity spatial interpolation for cellular coverage analysis, in: Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt), 2014 12th International Symposium on, IEEE, 2014, pp. 188–195.
- [21] T. Dasu, D. F. Swayne, D. Poole, Grouping Multivariate Time Series: A Case Study, in: Proceedings of the IEEE Workshop on Temporal Data Mining: Algorithms, Theory and Applications, in conjunction with the Conference on Data Mining, Houston, 2005, pp. 25–32.
- [22] A. Strehl, J. Ghosh, R. Mooney, Impact of similarity measures on webpage clustering, in: Workshop on Artificial Intelligence for Web Search (AAAI 2000), 2000, pp. 58–64.
- [23] MobiPerf, http://www.mobiperf.com(2014).
- [24] A. Nikravesh, D. R. Choffnes, E. Katz-Bassett, Z. M. Mao, M. Welsh, Mobile Network Performance from User Devices: A Longitudinal, Multidimensional Analysis, in: Procs. of PAM, 2014.
- [25] J. Sommers, P. Barford, Cell vs. WiFi: On the Performance of Metro Area Mobile Connections, in: Proc. of IMC, 2012.
- [26] J. P. Rula, V. Navda, F. Bustamante, R. Bhagwan, S. Guha, "No One-Size Fits All": Towards a principled approach for incentives in mobile crowdsourcing, in: Proc. of IMC, 2014.
- [27] M. Z. Shafiq, L. Ji, A. X. Liu, J. Wang, Characterizing and Modeling Internet Traffic Dynamics of Cellular Devices, in: Proc. of SIGMETRICS, 2011.
- [28] E. Halepovic, J. Pang, O. Spatscheck, Can you GET me now?: Estimating

the time-to-first-byte of HTTP transactions with passive measurements., in: Proc. of IMC, 2012.

- [29] M. Z. Shafiq, L. Ji, A. X. Liu, J. Pang, S. Venkataraman, J. Wang, A first look at cellular network performance during crowded events, in: Proc. of SIGMETRICS, 2013.
- [30] J. Huang, F. Qian, Y. Guo, Y. Zhou, Q. Xu, Z. M. Mao, S. Sen, O. Spatscheck, An In-depth Study of LTE: Effect of Network Protocol and Application Behavior on Performance, in: Proc. of SIGCOMM, 2013.
- [31] P. Benko, G. Malicsko, A. Veres, A large-scale, passive analysis of end-toend tcp performance over gprs, in: INFOCOM 2004. Twenty-third AnnualJoint Conference of the IEEE Computer and Communications Societies, Vol. 3, IEEE, 2004, pp. 1882–1892.
- [32] F. Ricciato, P. Svoboda, J. Motz, W. Fleischer, M. Sedlak, M. Karner, R. Pilz, P. Romirer-Maierhofer, E. Hasenleithner, W. Jäger, et al., Traffic monitoring and analysis in 3g networks: lessons learned from the metawin project, e & i Elektrotechnik und Informationstechnik 123 (7-8) (2006) 288–296.
- [33] S. Sen, J. Yoon, J. Hare, J. Ormont, 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.
- [34] Z. Koradia, G. Mannava, A. Raman, G. Aggarwal, V. Ribeiro, A. Seth, S. Ardon, A. Mahanti, S. Triukose, First Impressions on the State of Cellular Data Connectivity in India, in: Procs. of ACM DEV-4, ACM DEV-4 '13, 2013.
- [35] D. Baltrūnas, A. Elmokashfi, A. Kvalbein, Measuring the Reliability of Mobile Broadband Networks, in: Proc. of IMC, 2014.
- [36] M. Laner, P. Svoboda, P. Romirer-Maierhofer, N. Nikaein, F. Ricciato, M. Rupp, A comparison between one-way delays in operating hspa and lte networks, in: Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt), 2012 10th International Symposium on, IEEE, 2012, pp. 286–292.
- [37] D. Baltrūnas, A. Elmokashfi, A. Kvalbein, Dissecting Packet Loss in Mobile Broadband Networks from the Edge, in: Proc. of INFOCOM, 2015.
- [38] P. Perala, A. Barbuzzi, G. Boggia, K. Pentikousis, Theory and Practice of RRC State Transitions in UMTS Networks, in: IEEE GLOBECOM Workshops, 2009.
- [39] F. Qian, Z. Wang, A. Gerber, Z. M. Mao, S. Sen, O. Spatscheck, Characterizing Radio Resource Allocation for 3G Networks, in: Proc. of IMC, 2010.
- [40] S. Rosen, H. Luo, Q. A. Chen, Z. M. Mao, J. Hui, A. Drake, K. Lau, Discovering Fine-grained RRC State Dynamics and Performance Impacts in Cellular Networks, in: Proc. of Mobicom, 2014.
- [41] Y. Chen, N. Duffield, P. Haffner, W. ling Hsu, G. Jacobson, Y. Jin, S. Sen, S. Venkataraman, Z. li Zhang, Understanding the Complexity of 3G UMTS Network Performance, in: IFIP Networking Conference, 2013.
- [42] A. Gember, A. Akella, J. Pang, A. Varshavsky, R. Caceres, Obtaining In-Context Measurements of Cellular Network Performance, in: Proc. of IMC, 2012.
- [43] L. Li, K. Xu, D. Wang, C. Peng, Q. Xiao, R. Mijumbi, A Measurement Study on TCP Behaviors in HSPA+ Networks on High-speed Rails, in: Proc. of INFOCOM, 2015.
- [44] A. Balachandran, V. Aggarwal, E. Halepovic, J. Pang, S. Seshan, S. Venkataraman, H. Yan, Modeling Web Quality-of-Experience on Cellular Networks, in: Proc. of MobiCom, 2014.
- [45] F. P. Tso, J. Teng, W. Jia, D. Xuan, Mobility: A Double-Edged Sword for HSPA Networks, in: Proc. of MobiHoc, 2010.
- [46] P. Svoboda, F. Ricciato, W. Keim, M. Rupp, Measured web performance in gprs, edge, umts and hsdpa with and without caching, in: World of Wireless, Mobile and Multimedia Networks, 2007. WoWMoM 2007. IEEE International Symposium on a, IEEE, 2007, pp. 1–6.
- [47] A. Botta, A. Pescapé, Monitoring and measuring wireless network performance in the presence of middleboxes, in: Wireless On-Demand Network Systems and Services (WONS), 2011 Eighth International Conference on, IEEE, 2011, pp. 146–149.
- [48] V. Farkas, B. Héder, S. Nováczki, A Split Connection TCP Proxy in LTE Networks, in: Information and Communication Technologies, Springer, 2012, pp. 263–274.
- [49] X. Xu, Y. Jiang, T. Flach, E. Katz-Bassett, D. Choffnes, R. Govindan, Investigating Transparent Web Proxies in Cellular Networks, in: Proc. of Passive and Active Measurement, 2015.
- [50] J. Hui, K. Lau, A. Jain, A. Terzis, J. Smith, How YouTube Performance

is Improved in T-Mobile Network, in: Proc. of Velocity, 2014.[51] Z. Wang, Z. Qian, Q. Xu, Z. Mao, M. Zhang, An untold story of middle-boxes in cellular networks, in: Proc. of SIGCOMM, 2011.