

# Commute Path Bandwidth Traces from 3G Networks: Analysis and Applications

Haakon Riiser<sup>1</sup>, Paul Vigmostad<sup>1</sup>, Carsten Griwodz<sup>2,3</sup>, Pål Halvorsen<sup>2,3</sup>

<sup>1</sup>Opera Software ASA, Norway

<sup>2</sup>Simula Research Laboratory, Norway

<sup>3</sup>Department of Informatics, University of Oslo, Norway

{haakon.riiser, paulv}@opera.com  
{griff, paalh}@ifi.uio.no

## ABSTRACT

In this dataset paper, we present and make available real-world measurements of the throughput that was achieved at the application layer when adaptive HTTP streaming was performed over 3G networks using mobile devices. For the streaming sessions, we used popular commute routes in and around Oslo (Norway) traveling with different types of public transportation (metro, tram, train, bus and ferry). We also have a few logs using a car. Each log provides a timestamp, GPS coordinates and the measured number of bytes downloaded for approximately every second of the route. The dataset can be used in several ways, but the most obvious application is to emulate the same network bandwidth behavior (on specific geographical positions) for repeated experiments.

## Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: [Wireless communication]

## General Terms

Experimentation, Measurement, Performance

## Keywords

bandwidth traces, adaptive streaming, bitrate adaption, mobile internet, 3G, wireless, fluctuating bandwidth

## 1. INTRODUCTION

Mobile video streaming is an active area of research and development that is driven both by the tremendous growth in smartphones and tablets, and by modern wireless networks capable of streaming video in real-time even while the receiver is moving. For example, YouTube reports [4] that

*“traffic from mobile devices tripled in 2011, ... more than 20 % of global YouTube views come from mobile devices, and ... YouTube is available on 350 million devices.”*

Similarly, Sandvine reports that “real-time entertainment is huge, global and growing” [16, 15] for mobile devices, where North America, Latin America, Europe and Asia-Pacific have audio/video downstream mobile traffic of 27 %, 24 %, 17 % and 14 %, respectively. Sandvine also predicts that audio and video streaming will exceed 60 % of North America’s mobile data by late 2014. For mobile devices, Cisco’s Visual Networking Index predicts an 18-fold increase from 2011 to 2016 [2]. This trend is also enabled and fueled by all major content and service providers that make video delivery “apps” to popular mobile video platforms. Thus, the amount of media data streamed to mobile devices comprises already a large part of the total data downloaded in the Internet, and the scope and size of mobile video streaming is increasing at a very fast pace.

One challenge when delivering time-dependent data such as audio and video to moving devices is the changing network conditions and the resulting fluctuations in available bandwidth. HTTP-based adaptive streaming [17], which is tolerant to these fluctuations, is rapidly adopted by commercial vendors of streaming technology to be able to adapt to resource availability. Using adaptive streaming protocols such as Microsoft’s Smooth Streaming [18], Apple’s HTTP Live Streaming [9] or MPEG’s Dynamic Adaptive Streaming over HTTP (DASH), adaptive HTTP streaming is now also available on mobile devices. An adaptive streaming service can be implemented as a combination of simple servers and intelligent clients that make adaptation decisions based on local observations such as the observed bandwidth availability.

An important question in these scenarios is how to best adapt the video quality (and thus the bitrate) to the available network bandwidth. Some streaming systems available today have been designed for wireless streaming, but delivering video without interruptions is considerably more challenging when the client’s access network is a mobile wireless network with severe and frequent bandwidth fluctuations and outages. In such scenarios, the behavior of the quality adaption schemes varies greatly among the different solutions [8, 10]. To avoid recurring buffer underruns and frequent quality switches, both of which are harmful to the viewer’s quality of experience, several research groups have

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MMSys’13, February 26 - March 1, 2013, Oslo, Norway.

Copyright 2013 ACM 978-1-4503-1894-5/13/02 ...\$15.00.

evaluated the performance of adaption schemes [5, 8, 10] (and proposed changes [11]). However, they all use different bandwidth traces due to lack of an available dataset, e.g., Akhshabi et al. [5] used a synthetic bandwidth, Müller et al. [8] used real-world bandwidth traces from Klagenfurt (Austria) and Riiser et al. [10] used real-world bandwidth traces from Oslo (Norway). Even though the experiments show similar trends [8, 10], they are not directly comparable since the authors have used different network traces. With the dataset made available with this paper, we allow future researchers to reuse our traces, which we created during investigations resulting in several articles [6, 7, 10, 11, 13], and compare the results directly.

The dataset mostly contains bandwidth logs from Telenor’s 3G mobile network in and around Oslo. The network is based on the Universal Mobile Telecommunications System (UMTS) and the High-Speed Downlink Packet Access (HSDPA) architectures. The dataset includes bandwidth measurements from popular commute routes in Oslo and a few other examples. The commute routes all have multiple samples, and all routes come with a map of the path. Furthermore, we present statistics from the network traffic. This data was collected at the application-level in a HTTP-based media streaming client. Hence, it does not contain packet-level information similar to tcpdump. Bandwidth numbers are one-second averages, as this was sufficient for the media streaming simulations for which the data was originally intended. We hope that the dataset can be used by others to advance research in the field and provide better system support for mobile applications like adaptive video streaming.

## 2. MEASUREMENTS AND LOGS

We have earlier reported significant fluctuations in network conditions when streaming video over HSDPA networks [12]. To summarize our earlier results, we have performed bandwidth measurements in real-world streaming scenarios along several commute routes in and around Oslo, Norway. Based on these measurements, we have built a dataset over multiple routes by logging the download bandwidth observed while streaming video over HTTP. As a mobile video receiver, we used the Netview Media Client<sup>1</sup> [1] running on laptops equipped with a GPS (Haicom HI-204III USB GPS) and a 3G modem (Huawei Model E1752 HSPA USB stick). The sender was a dedicated server with a 100 Mbit/s Ethernet connection, located near the receiver. This was done so that our bandwidth and latency measurements indicate the performance of the wireless 3G connection, with minimal interference from the wired network. For example, we observed packet round-trip times of approximately 80 ms between the server and the client, and the round-trip time between the server and the last node before the wireless hop was less than 2 ms. Thus, more than 97 % of the observed end-to-end latency was due to the final wireless hop.

The Netview Media Client uses GPSD [3] to communicate with the GPS device, and counts the number of bytes received every second. It periodically reports to a remote database the GPS coordinates, timestamps, and bytes received since the last measurement. The dataset (logs, maps,

<sup>1</sup>The Netview Media Client that was used to record this data is proprietary (owned by Opera Software), hence we cannot distribute it together with the dataset.

|            |        |           |           |        |      |
|------------|--------|-----------|-----------|--------|------|
| 1289406399 | 549692 | 59.851754 | 10.781778 | 248069 | 1008 |
| 1289406400 | 550772 | 59.851864 | 10.781833 | 191698 | 1080 |
| 1289406401 | 551773 | 59.851964 | 10.781901 | 280579 | 1001 |
| 1289406402 | 552893 | 59.85206  | 10.781969 | 248971 | 1120 |

Figure 1: Sample log data.

and bandwidth plots) and custom-made software useful for simulations have been made available online [14]. The dataset currently consists of 86 traces from 11 different routes. The logs are in plain ASCII text, with six fields of data per log entry, as shown in figure 1. The data fields are defined as follows:

- Column 1 is the **Unix timestamp** (number of seconds since 1970-01-01) of the data point.
- Column 2 is a **monotonically increasing timestamp in milliseconds** (since some unspecified starting point).
- Columns 3 and 4 are **GPS coordinates in decimal degrees**. Column 3 is the latitude coordinate and column 4 is the longitude coordinate.
- Column 5 is the **number of bytes received** since the previous measurement (the previous line in the log).
- Column 6 is the **number of milliseconds elapsed** since the previous measurement (the previous line in the log). Column 6 equals the difference in the column 2 values of this and the previous line.

This means that other types of information can easily be extracted, e.g., to get kilobytes per second for a specific sample, simply divide column 5 by column 6.

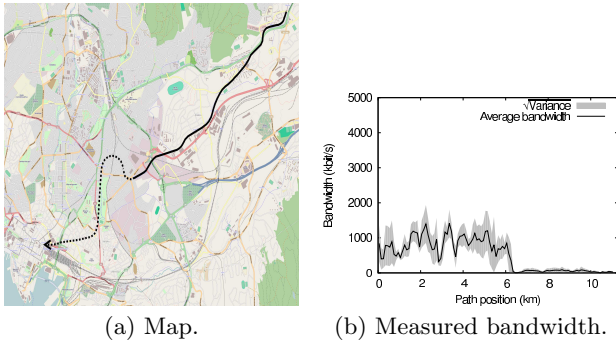
## 3. ROUTES AND BANDWIDTHS

In the previous section, we presented the logs, and here, we give some examples of the used paths with corresponding bandwidth measurement information. With a few exceptions, we selected popular commute routes to or from downtown Oslo, Norway. For different routes, we also used different types of public transportation (metro, bus, ferry, train, tram and car). In the following subsections, we list a few examples (of the total 11 routes) with maps of the routes and plots of the average observed bandwidth over multiple measurements as a function of the path position (as traveled distance from the start). Note that some of the maps are slightly stretched to fit the format, and that the measured bandwidths in the bandwidth plots are calculated over a one second average.

### 3.1 Metro railway

A popular means of commuting in Oslo is the metro. This is an electric passenger railway where all the lines pass through underground tunnels in downtown Oslo. The underground part of our tested metro commute route is shown with the dotted line in figure 2(a).

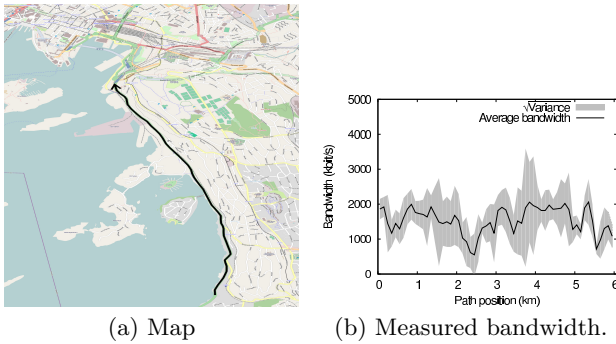
Figure 2(b) shows the measured bandwidth along the metro path. All the measurements show the same trend, and the signal and bandwidth availability are predictable with only minor variations. The experienced bandwidth is typically a bit above 1 Mbit/s when the metro is above the ground.



**Figure 2: Metro commute path: from Kalbakken to Jernbanetorget in Oslo (the dotted line is in a tunnel).**

However, when entering the tunnels after approximately 5.5 kilometers, both the Internet connection and the GPS signal are essentially lost (in periods without a GPS signal, the position is estimated based on the metro time table). This means that when the metro was underground (the dotted line), we were hardly able to receive any data at all.

### 3.2 Bus

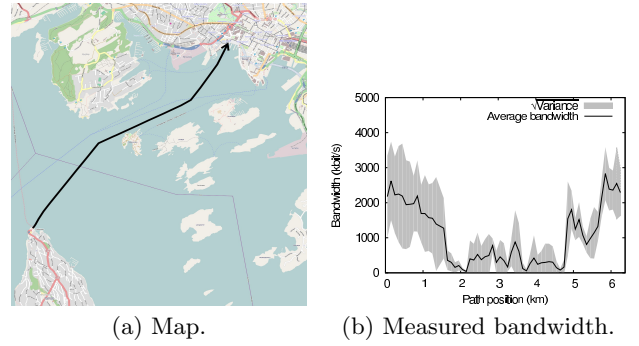


**Figure 3: Bus commute path: Oslo bus along Mosseveien, between Ljan and Oslo Central Station.**

Figures 3(a) and 3(b) show a bus path going into Oslo and the corresponding bandwidth measurements, respectively. The average values in figure 3(b) vary greatly, but the measurements show that a minimum bandwidth of about 1.5 Mbit/s should normally be possible with an average of about 2 Mbit/s. A challenge using this route is that the bus route has a steep hill on the east side, which prevents a stable reception of signals from eastern GPS satellites. Furthermore, the Oslofjord is in the west, leaving few possible sites for 3G towers on that side. Consequently, both the GPS and the UMTS signals are unstable.

### 3.3 Ferry

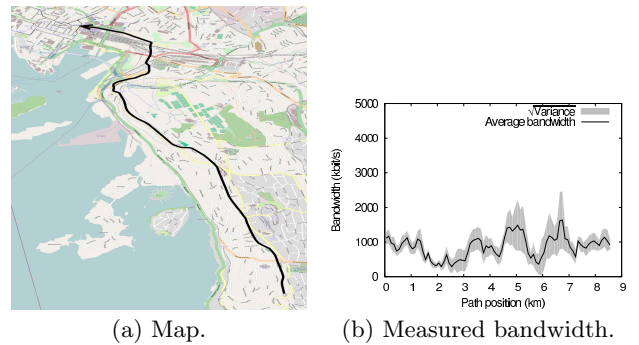
Our third scenario is traveling by ferry (figure 4(a)). Most commuters from the Nesodden peninsula travel by ferry to Oslo, as traveling by car requires a large detour. With lots



**Figure 4: Ferry commute path: Oslo ferry route between Nesodden and Aker Brygge.**

of space on board and tables for PCs and devices with large screens, this way of commuting is one that is well suited to use mobile devices during the ferry ride. Figure 4(b) shows that the available bandwidth depends strongly the position along the trip. The signal is strongest when the ferry is close to land (Nesodden at the start of the path, and Aker Brygge in downtown Oslo) with observed average bandwidths above 2 Mbit/s. However, the signal conditions far from land are usually problematic. The signal is never completely gone while crossing the Oslo-fjord, but the 3G users rarely experience bandwidths above 1 Mbit/s in this part of the path.

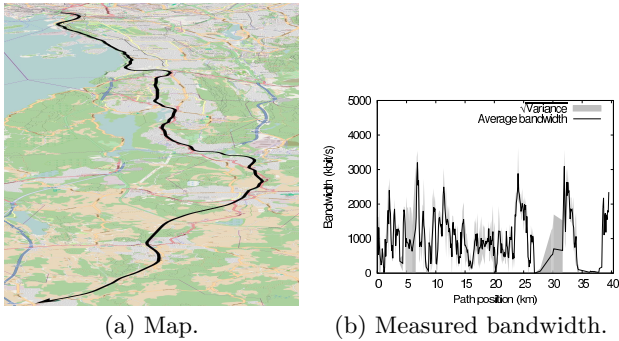
### 3.4 Tram



**Figure 5: Tram commute path: Oslo tram route between Ljabru and Oslo Central Station.**

Another method of commuting in Oslo using public transportation is by tram. Figure 5(a) shows one of the tested tram routes, whose tracks are parallel to but high above the bus route presented above. Observed bandwidth is presented in figure 5(b). Along the whole path, we have found acceptable but fluctuating bandwidth. In the first part of the route, we observed a very predictable bandwidth across the different measurements, though rarely exceeding 1 Mbit/s. At a long down-slope towards the end of the trip, the measurements vary more (larger standard deviation), but the average bandwidth is almost the same.

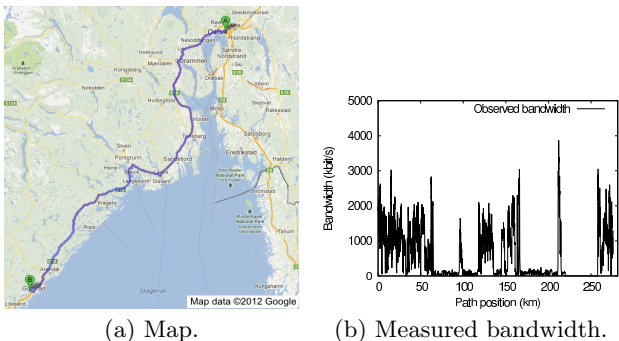
### 3.5 Train



**Figure 6: Train commute path: Train route between Oslo and Vestby.**

The trains to and from Oslo are frequently used by people traveling longer distances, and in figure 6, we show the map and bandwidth plot for the 40 km long Oslo–Vestby route. We observe large fluctuations in bandwidth, sometimes jumping between 3 Mbit/s and almost no connectivity at all.

### 3.6 Car



**Figure 7: Car path: Driving a car from Oslo to Grimstad.**

The route Oslo–Grimstad, shown in figure 7(a), is used by everyone driving from Oslo going south on the E18 highway. It is an approximately 280 kilometer drive. Figure 7(b) shows the achieved bandwidth, where we observed high peaks over 3 Mbit/s with an average of about 1 Mbit/s. However, as also seen in the plot, there are several areas with (nearly) no available bandwidth.

### 3.7 Summary

In summary, the above plots indicate that the achieved bandwidth oscillates severely when using moving mobile devices. How much depends on the geographical location and the speed of the moving vehicle, and this behavior is definitely something that should be taken into account when designing systems for the mobile scenario.

## 4. EXAMPLES OF USE

In our work, we have used the bandwidth measurements to emulate identical network behavior for video streaming sessions using adaptive HTTP streaming protocols. The goal was to evaluate different adaption strategies, i.e., decisions on how to change from one quality (and thus bitrate) level to another, as a user moves along a path while viewing a video in a Video-on-Demand scenario [10, 11]. Additionally, we have used the logs to emulate the network in a scenario where the logs are used to make a location-based bandwidth-lookup service [13, 11] which also may be combined with a transparent handover between multiple networks [7, 6].

To produce the same network conditions for all tests, we developed a bandwidth throttling module for the Apache web server (also available from the dataset web-page). This module takes as input a bandwidth log, like the logs available in this dataset, that contains a single kbit/s number for every second of the session. After loading the bandwidth log, the first HTTP request starts the session. At time  $t$  after the session starts, the web server’s maximum throughput for the next second will be  $B(t)$ , where  $B(t)$  is the bandwidth at time  $t$  in the log that was used as input to the throttling module. This approach means that the client program, i.e., the media player in our HTTP streaming scenario, can get exactly the same bandwidth conditions every time, ensuring both fairness and reproducibility in our experiments, while at the same time being as realistic as a field trial.

## 5. CONCLUSION

We have provided a dataset that enables researchers to simulate bandwidth fluctuations as observed in a live 3G network in vehicular mobility scenarios. Using this dataset, researchers can test their ideas using realistic network traces, and get reproducible results useful for comparisons with other researchers. We hope that the dataset can be used to push research forward in the field of data delivery to mobile devices, and in streaming scenarios in particular.

## Acknowledgements

This work has been performed in the context of the *HyStream* project (project number 176847) and the *iAD* centre for Research-based Innovation (project number 174867) – both funded by Norwegian Research Council.

## 6. REFERENCES

- [1] Netview Media Client. <http://www.netview.no/index.php?page=downloader>, 2011.
- [2] Cisco visual networking index: Forecast and methodology, 2011-2016. [http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white\\_paper\\_c11-481360\\_ns827\\_Networking\\_Solutions\\_White\\_Paper.html](http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360_ns827_Networking_Solutions_White_Paper.html), 2012.
- [3] GPSD — a GPS service daemon. <http://catb.org/gpsd/index.html>, 2012.
- [4] YouTube statistics. [http://www.youtube.com/t/press\\_statistics](http://www.youtube.com/t/press_statistics), Nov. 2012.

- [5] S. Akhshabi, A. C. Begen, and C. Dovrolis. An experimental evaluation of rate-adaptation algorithms in adaptive streaming over http. In *Proceedings of the ACM International Conference on Multimedia Systems (MMSys)*, pages 157–168, Feb. 2011.
- [6] K. Evensen, A. Petlund, H. Riiser, P. Vigmostad, D. Kaspar, C. Griwodz, and P. Halvorsen. Demo: Quality-adaptive video streaming with dynamic bandwidth aggregation on roaming, multi-homed clients (demo). In *Proceedings of the ACM International Conference on Mobile Systems, Applications and Services (MobiSys)*, June 2011.
- [7] K. Evensen, A. Petlund, H. Riiser, P. Vigmostad, D. Kaspar, C. Griwodz, and P. Halvorsen. Mobile video streaming using location-based network prediction and transparent handover. In *Proceedings of the ACM International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV)*, June 2011.
- [8] C. Müller, S. Lederer, and C. Timmerer. An evaluation of dynamic adaptive streaming over http in vehicular environments. In *Proceeding of the ACM Workshop on Mobile Video (MoVid)*, Feb. 2012.
- [9] R. Pantos and W. May. HTTP live streaming. <http://tools.ietf.org/html/draft-pantos-http-live-streaming-07>, 2011.
- [10] H. Riiser, H. S. Bergsaker, P. Vigmostad, C. Griwodz, and P. Halvorsen. A comparison of quality scheduling in commercial adaptive http streaming solutions on a 3g network. In *Proceedings of the ACM Workshop on Mobile Video (MoVid)*, Feb. 2012.
- [11] H. Riiser, T. Endestad, P. Vigmostad, C. Griwodz, and P. Halvorsen. Video streaming using a location-based bandwidth-lookup service for bitrate planning. *ACM Transactions on Multimedia Computing, Communications and Applications (TOMCCAP)*, 8(3), 2012.
- [12] H. Riiser, P. Halvorsen, C. Griwodz, and B. Hestnes. Performance measurements and evaluation of video streaming in HSDPA networks with 16QAM modulation. In *Proceedings of the IEEE International conference on Multimedia and Expo (ICME)*, pages 489–492, June 2008.
- [13] H. Riiser, P. Vigmostad, C. Griwodz, and P. Halvorsen. Bitrate and video quality planning for mobile streaming scenarios using a gps-based bandwidth lookup service. In *Proceedings of the IEEE International Conference on Multimedia and Expo (ICME)*, July 2011.
- [14] H. Riiser, P. Vigmostad, C. Griwodz, and P. Halvorsen. DATASET: HSDPA-bandwidth logs for mobile HTTP streaming scenarios. <http://home.ifi.uio.no/paalh/dataset/hsdpa-tcp-logs/>, Oct. 2012.
- [15] Sandvine Intelligent Broadband Networks. Global internet phenomena report: 1h 2012. [http://www.sandvine.com/news/global\\_broadband\\_trends.asp](http://www.sandvine.com/news/global_broadband_trends.asp), Apr. 2012.
- [16] Sandvine Intelligent Broadband Networks. Sandvine report: Mobile networks teeming with streaming. [http://www.sandvine.com/news/pr\\_detail.asp?ID=366](http://www.sandvine.com/news/pr_detail.asp?ID=366), Apr. 2012.
- [17] T. Stockhammer. Dynamic adaptive streaming over HTTP: Standards and design principles. In *Proceedings of the ACM International Conference on Multimedia Systems (MMSys)*, pages 133–144, Feb. 2011.
- [18] A. Zambelli. Smooth streaming technical overview. <http://learn.iis.net/page.aspx/626/smooth-streaming-technical-overview/>, 2009.