Open Video Datasets
over Operational Mobile Networks with MONROE

Cise Midoglu[1,4], Mohamed Moulay[2,3],Vincenzo Mancuso[2],
Özgü Alay[1,4], Andra Lutu[1], Carsten Griwodz[4,1]

[1] Simula Research Laboratory, Norway   [2] IMDEA Networks Institute, Spain
[3] University Carlos III of Madrid, Spain   [4] University of Oslo, Norway

ABSTRACT

Video is becoming a killer application for mobile users. DASH and WebRTC are the brightest stars in this context. To assess the performance of DASH and WebRTC, we use a large number of programmable network probes spread over several countries in the context of the MONROE project. The probes allow to analyze the end-user performance in a repeatable and controllable way over operational mobile networks. We have generated a large dataset with more than 300 video streaming experiments so far, and growing. Our dataset consists of traces from video streams, performance indicators captured during the streaming, and experiment metadata. The dataset captures the wide variability of video quality and unveils how broadband mobile access is still not offering consistent quality guarantees across several monitored countries and networks, especially for users on the move. We make both software and datasets (video data and metadata) publicly open.

1 INTRODUCTION

Mobile video traffic accounted for 60% of total mobile data traffic in 2016, and by 2021, over 75% of the world's mobile data traffic will be video [2]. Dynamic Adaptive Streaming over HTTP (DASH) and WebRTC are two key techniques for the delivery of mobile video in a standardized manner. DASH is used for adaptive on-demand streaming, and DASH-compliant applications provide users with various network adaptation algorithms [10]. WebRTC is a solution that enables web browsers to communicate with each other and supports voice, video and data communication in a peer-to-peer fashion, in principle with no need of servers, plugins and additional software [22]. The popularity of DASH and WebRTC continues to grow due to the increased availability of bandwidth offered by the Mobile Network Operators (MNOs). This makes it possible to deliver a good service, although users of video services report unreliable quality of experience [21]. Therefore, understanding how DASH and WebRTC work and under which conditions they do a good job is crucial for designing services with better and certified quality of experience guarantees, especially for the constantly increasing number of mobile customers [9].

In this context, an empirical analysis of video delivery over mobile networks is paramount because new services are continuously proposed and do not wait for approximated models and network optimization studies before going live. Indeed, they go to the market and are sold without waiting for MNOs to adjust [1]. However, it is notoriously difficult to obtain systematic data from customers through repeatable measurements and, at the same time, collect proper information on the conditions under which these measurements run [11, 19, 20]. To overcome this problem, we have designed, developed and deployed a dedicated hardware platform to monitor the behavior of Mobile Broadband (MBB) networks and the applications running on them, in the context of the Measuring Mobile Broadband Networks in Europe (MONROE) project [3]. MONROE runs measurement probes from a user perspective, i.e., from the access network and by relying on regular mobile data subscriptions. MONROE nodes are currently operational in seven European countries, from where they can monitor more than 15 MNOs.

In this paper, we focus on Video on Demand (VoD) streaming and we empirically assess mobile streaming performance by leveraging MONROE nodes to run streaming experiments from stationary and mobile measurement vantage points. We evaluate performance of different rate adaptation algorithms with DASH. To evaluate VoD with Real-time Transport Protocol (RTP), we use WebRTC in a server-to-client video streaming mode. Our dataset is continuously enriched with fresh data and captures the variable end-user video experience *in the wild*, in different countries and for different MNOs. It also enables continuous monitoring of how services and networks evolve over time, and the comparison of different experiments in time and space. The dataset includes DASH and WebRTC traces and performance indices jointly with a generous quantity of *metadata*. Our dataset currently has results from over 300 individual experiments with log files and more than 500K lines of metadata for GPS, sensor and modem information. This is possible because MONROE nodes are not only instrumented to periodically measure the performance of their MNOs, but also store node location, signal strength, cellular technology used, etc.

The measurements run as Docker containers in which we have packaged our measurement tools, and which we distribute as open source [14]. Specifically, here we focus on two containers for quantifying DASH and WebRTC performance, namely `AStream` [6] and `WebStreamer` [18]. The former implements DASH with 3 alternative adaptive algorithms, whereas the latter implements WebRTC using HTTPS. We show that our containers enable monitoring and assessing DASH and WebRTC performance over mobile broadband connections from stationary and mobile devices in a repeatable and controllable way. Furthermore, we describe the datasets resulting from the fusion of measurement data and metadata and make our measurements publicly available as an open video dataset [8].

2 DATA COLLECTION

In this section, we describe the way the MONROE video dataset is collected.

2.1 MONROE Platform

MONROE [3, 17] is a European transnational open platform, and the first open access hardware-based platform for independent, multi-homed, and large-scale MBB measurements on commercial networks. The platform comprises a set of 150 nodes, both mobile (e.g., operating in delivery trucks and on board public transport vehicles, such as trains or buses) and stationary (e.g., volunteers hosting nodes in their homes). Each node is multi-homed to 3 different MNOs. MONROE is currently operational in Italy, Norway, Spain, Sweden, Portugal, Greece and the UK. All software components used in the platform are open source and available online [14].

User access to the platform resources is through a web portal, which allow authenticated users to use the MONROE scheduler to deploy their experiments. This enables exclusive access to nodes (i.e., no two experiments run on a node at the same time). Experiments running on the platform use Docker containers (light-weight virtualized environment) to provide agile reconfiguration.

The designed DASH and WebRTC experiments can be run easily on MONROE with different configuration parameters, and measurement results from operational MBB networks can be obtained in real time. We provide templates with default parameters to run the containers directly on the MONROE scheduler interface, so that both experiments can be run with a single click.

Metadata collection. Since MONROE does not involve real users (which usually entail privacy protection restrictions), rich metadata collection, including geo-temporal tagging, is possible. MONROE nodes generate metadata passively and continuously: each node is instrumented to gather information relating to its MNOs, such as location, signal strength, and link technology.

At the node side, metadata distribution is implemented in a publish/subscribe pattern using ZeroMQ (ZMQ).[1] Metadata entries are generated in a single-line JavaScript Object Notation (JSON) format, where every entry is labeled with a "topic" field. The metadata subscriber module subscribes to all the topics, writing JSON entries to files in a special file system location. A synchronization process transfers these files to the MONROE server when no other active, periodic, or user-defined experiment is running. In this way, metadata from all MONROE nodes is collected and stored centrally.

The metadata stream is also available for experiments during their execution. Experimenters can have their containers subscribe to any of the metadata topics, and use them either during runtime or store them together with experiment results. The former allows for monitoring and reacting to events such as interface reconnections or link technology/signal strength change for each MBB interface at run-time, where the latter facilitates the joint post-processing of measurement data and metadata pertinent to a given experiment.

2.2 DASH Experiments

We design the MONROE-AStream experiment for conducting DASH measurements over operational MNOs using the MONROE platform. The overall framework consists of 5 steps and is described in Figure 1 as a flow diagram.

First, a video is selected for streaming. This video can be from any source (e.g. YouTube). The DASH representations of the video are downloaded/generated and the retrieved segments are used

[1]ZMQ distributed messaging: http://zeromq.org

Figure 1: MONROE-AStream framework.

for Media Presentation Description (MPD) generation. An existing MPD file can also be used.

Then, the MPD file and the segments are uploaded to the server. The original AStream server script is modified to accommodate for the new MPD file and the location of different video representations. The server is configured to listen to a desired port. It is also possible to use an existing DASH server, which is able to respond to client requests regarding the selected video, and compatible with its corresponding MPD file.

The MONROE-AStream client consists of the AStream core with an additional patch for reading different MPD formats [4], configuration and logging functionality for MONROE-related parameters, wrapped in a MONROE-compatible Python script. The client container and the server code can run both within or outside of the MONROE platform, allowing for measurements in different setups.

The final step of our experiment framework is the analysis of output files, for which our toolset includes parsing and plotting scripts. The experiment container is provided as open source within an Experiment as a Service (EaaS) framework, under [6].

Figure 2 depicts our experiment scenario where the MONROE platform is used as the measurement environment in Step 4. We aim to investigate the performance of different rate adaptation algorithms over different MNOs. The rate adaptation algorithms we consider are: Basic, Segment Aware Rate Adaptation (SARA) [13], and Buffer Based Adaptation (BBA) [12].


Figure 2: MONROE-AStream measurement setup.

2.3 WebRTC Experiments

The main idea for the MONROE-WebRTC experiment is to test video streaming using WebRTC for mobile users. Figure 3 outlines the streaming setup. When an experiment is scheduled on a set of MONROE nodes, each node runs the WebStreamer container, which makes an HTTPS link available to watch the video stream coming from the node. We use a Chrome browser acting as the WebRTC client in our lab. The WebRTC client is connected from a computer wired to a well provisioned gigabit link so that the receiver is not
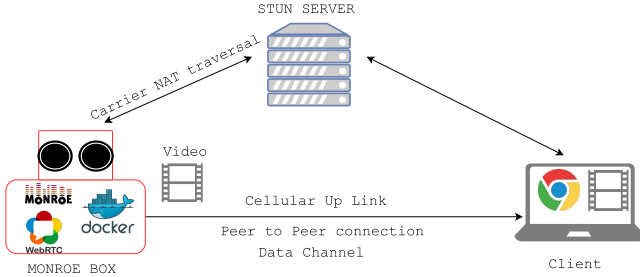
Figure 3: MONROE-WebRTC video streaming setup.

the bottleneck. The video stream produced by the MONROE node goes through a cellular uplink, traverses the Internet and then accesses the network of our lab. Therefore, the network bottleneck experienced in a WebRTC streaming session is the MBB network used by the MONROE node.

The video to stream is stored locally in the Docker container, with a resolution of 854x480 and a constant frame rate of 24 fps. The duration of the video is 9 minutes but can be looped for several hours, according to the duration of the scheduled experiment. The video streaming configuration parameters include the following:

- The possibility to use a local or online video by passing an RTSP link.
- Picking the preferred port for communication.
- Using a specific STUN server in case the client is behind a firewall.
- Tweaking resolution and duration of the video to stream.

At the client side, the video stream is received by means of Google Chrome, which also measures and collects streaming statistics in JSON format. The resulting dump contains Peer-to-Peer Connection status information between the MONROE node and the WebRTC client in our lab, in addition to updates and data statistics that can be easily accessed, e.g., by loading a specific Chrome page.[2]

The client logs contain, per each individual stream, the timing and headers of packets received as well as the timing of various internal events such as received frames, losses, bitrate, delay, jitter and other metrics that will be discussed in Section 3.

3 DATASET DETAILS

Our video dataset includes data and metadata relative to hundreds of experiments with stationary and mobile MONROE nodes in Italy, Norway, Sweden and Spain. This is a live dataset, so the dataset keeps growing. Experiment results included in the dataset so far have been collected within a period of eight months, and we are currently running more experiments to keep following the performance evolution in the monitored locations (for stationary nodes) and paths (for mobile MONROE nodes mounted on buses and trains). Our open dataset can be found under [8] and includes a README file explaining the organization of each directory (MONROE-AStream and MONROE-WebRTC).

Here, we provide a list of parameters that are available in the MONROE dataset relevant to video measurements. These come from the MONROE-AStream and WebStreamer experiments, and the MONROE metadata collection process.

---

[2]chrome://webrtc-internals

3.1 MONROE Metadata

Table 1 illustrates the metadata "topics", which are streamed to subscriber entities within MONROE nodes using ZMQ, as explained in Section 2, and the relevant experiments running regularly on the platform.

Table 1: MONROE Metadata topics and experiments.

Class	Type	Examples
Node	Sensor	CPU temperature
Node	Probe	Load, memory usage
Node	Event	Power up, reboot
Device	GPS	GPS coordinates
Device	Modem	RSSI, link technology, CID, IP
Experiment	RTT	Ping RTT
Experiment	Tstat	TCP statistics
Experiment	Bandwidth	Nettest TCP throughput

Metadata includes network parameters such as Received Signal Strength Indicator (RSSI), Reference Signal Received Power (RSRP), Reference Signal Received Quality (RSRQ), cell identifiers, link technology, node location and speed via GPS, node working parameters such as CPU temperature and load, and node events (watchdogs).

In our open dataset, we include samples from the `node.sensor` and `device.modem` classes for the duration of MONROE-AStream experiments, and samples from the `device.gps` class for the Web-Streamer experiments under mobility.

3.2 MONROE-AStream Dataset

MONROE-AStream produces an experiment summary in JSON format, 3 types of detailed logs, and the streamed segments in raw form (for which the download option is a configuration parameter). These are ready to download in the form of 2 files after experiment completion: one JSON file as the summary, and one compressed archive for all logs and video segments, both with the filename stub *MONROE.EXP.ASTREAM_<video>_ <adaptation>_<nodeid>_<interface>_<timestamp>*. Samples for each type of output can be found in our open dataset [8].

Experiment Summary. This is a JSON file with a single-level dictionary, generated per experiment. It lists the basic configuration parameters and platform-related information for the experiment, such as timestamp, node id, operator, container version, video name, adaptation algorithm, segment limit, and experiment tag.

Runtime Logs. These are `.log` files, providing the complete command line output of the code execution as a file, generated per experiment interface. They give a detailed overview of the experiment execution with timestamps, and enable troubleshooting.

Segment Logs. These logs are JSON files with a 4-level dictionary, including the following fields, generated per experiment interface.

- **playback_type:** indicates the selected adaptation algorithm (Basic, SARA, or BBA).
- **segment_info:** contains an entry per each streamed segment, indicating the segment filename, bitrate, size, and time to download.
- **playback_info:** contains the start time (*start_time*), end time (*end_time*), number of up switches (*up_shifts*), number

of down switches (*down_shifts*), duration of initial buffering in seconds (*initial_buffering_duration*), and information regarding buffering events, i.e. stalls (*interruptions*). The *interruptions* object includes the fields *count*, *total_duration*, and *events*.

- **video_metadata:** contains the total playback duration (*playback_duration*), and information regarding the available representations of the given video (*available_bitrates*). The latter includes an object with fields *width*, *bandwidth*, *height*, *frameRate*, and *codecs* per each representation.

Buffer Logs. These logs are Comma-Separated Values (CSV) files with event-based entries including the following fields, generated per experiment interface,

- **EpochTime:** unix timestamp.
- **CurrentPlaybackTime:** current playback time in seconds, with a granularity of segment duration (e.g. 2*s*).
- **CurrentBufferSize:** current buffer size in seconds.
- **CurrentPlaybackState:** current playback state as one of the following; INITIAL_BUFFERING, PLAY, STOP.
- **Action:** current action as one of the following; Starting, Writing, InitialBuffering-Play, StillPlaying, Stopped.
- **Bitrate:** current segment bitrate in bits per second.

Downloaded Segments. It is possible to download the streamed segments at the client side in raw form, with a configuration parameter. This option gives experimenters the ability to reconstruct the streamed video for playback.

3.3   MONROE-WebRTC Dataset

WebStreamer generates two types of log files. The first one is our Docker container log, which consists of the following:

- HTTP link to connect too.
  ```
  2018-xx-xxTxx:xx:xx your url is: https://xxxx
  ```
- The videos you are going to watch.
  ```
  2018-xx-xxTxx:xx:xx.xx answer:[
  2018-xx-xxTxx:xx:xx.xx "xx.mp4,rtsp://xxxxx"]
  ```
- The open port listening to upcoming connections.
  ```
  2018-xx-xxTxx:xx:xx.xx HTTP Listen at :8888
  ```
- The constant Framerate of the video.
  ```
  2018-xx-xxTxx:xx:xx.xx a=framerate:24.0
  ```
- The video codec used that includes h264, VP8, and VP9.
  ```
  2018-xx-xxTxx:xx:xx Created a data sink for the
                "video/H264/VP8/VP9" subsession
  ```
- The STUN server used when establishing connection between client and host.
  ```
  2018-xx-xxTxx:xx:xx.xx{"url" : "stun:stun.l.google.com:19302"
  2018-xx-xxTxx:xx:xx.xx        }
  ```
- The preferred candidate for connection.
  ```
  body:{"candidate":"candidate:0 2 UDP 2122252542
  192.xx.x.x 42479 typ host",
  "sdpMid":"sdparta_0","sdpMLineIndex":0}
  answer:1
  ```
- The Video Content Type which is an extension used to communicate a video content type from sender to receiver of RTP video stream.
  ```
  2018-xx-xxTxx:xx:xx.xx
  a=rtpmap:96 mpeg4-generic/48000/2
  ```

When the connection is established, the WebRTC receiver starts collecting the JSON dump from Google Chrome, which is our second log. This includes different metrics: *bitrate received, frames decoded,*

packets lost during the session, packets received per second, video current delay in ms, frame rate received, decoded and output, jitter delay in ms, number of ACKs sent, target delay in ms.

The MONROE-WebRTC template and the visualization tool that helps import all the metrics mentioned in CSV format can be found at [8] and [18].

4   APPLICATIONS OF THE DATASET

The MONROE video dataset is unique in two aspects: it incorporates video measurements from *operational* MBB networks, and it provides *context information* for all measurements. The former allows for the analysis of realistic MBB scenarios, including mobility, whereas the latter allows researchers to correlate their results with metadata. As an example, the performance of a rate adaptation algorithm may be related to the underlying network conditions during its execution, such as signal coverage.

MONROE-AStream and MONROE-WebRTC experiment results can be used for a number of purposes. For instance, it is possible to evaluate the performance of different MNOs, or the performance of different DASH rate adaptation algorithms [15]. It is also possible to compare performance under mobility with stationary scenarios [16]. In the following, we present an analysis regarding MNO and rate adaptation algorithm performance. Note that we do not report the name of the MNOs in our results because we do not intend to carry on an explicit comparison between MNOs.[3]

MONROE-AStream. We use the DASH server from ITEC with the readily available MPD file for the BigBuckBunny video from their open *MMSys* dataset (20 representations), selecting 2*s* duration for segments [5]. We run the MONROE-AStream client on stationary nodes in Norway and Sweden, using 6 operational MBB networks, with 3 rate adaptation algorithms (Basic, SARA, and BBA) with an option to download the streamed video segments. We specifically choose these countries since they have higher SIM quotas, allowing a larger number of measurements compared to other countries.

Using the MONROE-AStream experiments, it is possible to evaluate metrics such as segment bit rate, duration of stalls (buffering), number of up/down switches, and time spent at different video quality levels. It is possible to evaluate the performance of MNOs and adaptation algorithms separately, as well as compare their combined performance.

Figure 4 shows the boxplot of segment bitrate for each adaptation algorithm and operator in (a) Sweden and (b) Norway. As one of the most prominent indicators of video streaming quality, segment bitrate can be used to compare the performance of MNOs as well as adaptation algorithms. We see a higher overall video bitrate on average for nodes on Sweden when compared to nodes in Norway, and a higher bitrate on average for adaptation algorithms SARA and BBA compared to Basic.

Figure 5 shows the Empirical Cumulative Distribution Function (ECDF) of the total duration of up and down-switches for different MNOs, for a single adaptation algorithm (SARA). As an indicator of possible fluctuations in the available network throughput, as well as user experienced quality, the number of switches gives insights into

---

[3]The MNO names and exact GPS coordinates of the nodes used in the experiments are made available as metadata in the dataset.

**(a) Segment Bit Rate (Sweden)**   **(b) Segment Bit Rate (Norway)**

Figure 4: MONROE-AStream performance figures of segment bitrate from experiments run on stationary nodes in Sweden and Norway. Results are split according to operator and adaptation algorithm, per country.


(a) Up-Switches   **(b) Down-Switches**

Figure 5: MONROE-AStream performance figures as ECDF of number of quality switches per operator, from experiments run on stationary nodes in Sweden and Norway. Comparison of up and down-switches for SARA.

(a) op1   **(b) op2**

(c) op3   **(d) op4**
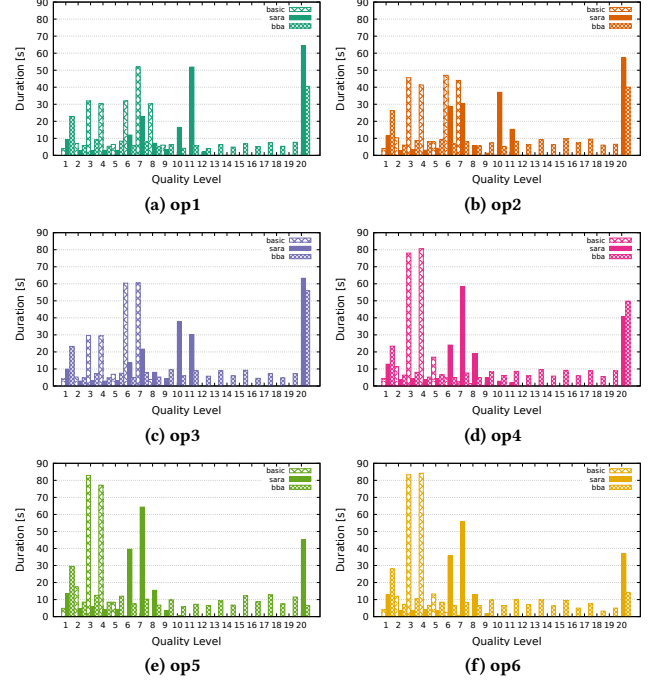
(e) op5   **(f) op6**

Figure 6: MONROE-AStream performance figures of average time spent at each video quality level per operator and adaptation algorithm, from experiments run on stationary nodes in Sweden and Norway. Results are split according to operator (op1-3 in Sweden, op4-6 in Norway).

the performance of video streaming as experienced by the clients. From our measurements, we observe that operators op2, op5, and op6 switch less in both directions as compared to operators op1, op3, and op4, on average.

Figure 6 shows the total time spent at each quality level for different operators. We observe for all operators that, while the Basic algorithm streams video segments mostly from the lower quality layers, SARA can stretch further to medium levels, with a jump to the highest quality after more than half of the total duration, and BBA traverses the low and medium quality levels quite quickly to spend most of the time at the highest quality level.

MONROE-WebRTC. Here, we use a sample of our dataset to demonstrate the performance analysis of WebRTC. We use WebRTC video data for a few stationary nodes across Sweden and Norway, and for three mobile nodes located in regional trains in Norway, in the Oslo metropolitan area and its surroundings. The video statistics we show here are the ones coming from the logs built by means of Chrome WebRTC internals, at the receiver side. We also use MONROE metadata generated at the MONROE nodes to complement and interpret the video data.

For experiments with stationary nodes, we select WebRTC streamers located aboard MONROE nodes with good MBB connectivity. Specifically, we consider WebRTC streams using as uplink an MBB connection under good 4G coverage. In Figure 7, we outline the average bitrate and average delay experienced during a connection. Error bars in the figure indicate the standard deviation observed in the statistics. A total number of 6 nodes were used, 6 different MNOs are measured. We split results according to the country of origin of the streams in the experiment, which is either Sweden or

Norway, whereas the receiver was located in all cases in Madrid, Spain. We further split the results per MNO.

As shown in the figure, the bitrate was consistently high. The figure also shows that the one-way delay observed during the video streaming sessions was acceptable and, in most of the cases, well below 150 ms, which is the upper limit beyond which a human being starts noticing the delay and gets annoyed.

For experiments with mobile streamers, we selected MONROE nodes traveling on trains. Here we report an example of data series obtained for three streams from three trains moving in and around Oslo in the same time interval. For this mobile case, Figure 8 depicts instantaneous values of bitrate received, video frame rate reconstructed at destination, and video delay and jitter experienced by the receiver. The figure shows that performance under mobility can be good, at least under good cellular coverage. However, having good cellular coverage does not seem to be always the case when moving in the observed area. By analyzing the metadata associated with the traces reported in Figure 8, we have noticed that one of the trains used for the experiment progressively moved from the city center to the outskirts of Oslo, while the other trains stayed within the city, moving between the multiple stations in Oslo. The streams from the trains remaining in the city experienced almost always good performance, although with high variability. On the other hand, the stream from the train moving outside of Oslo shows poor performance at the beginning of the experiment, when the stream suffered a huge degradation in bitrate and a significant increase in

**(a) Average bitrate**

(b) Average delay

Figure 7: MONROE-WebRTC average performance for the experiments run from stationary MONROE nodes in Sweden (op1-3) and Norway (op4-6). Results are split according to operator and country of origin of the WebRTC streams.

(a) Bitrate at the receiver

(b) Video delay

(c) Frames per second at the receiver
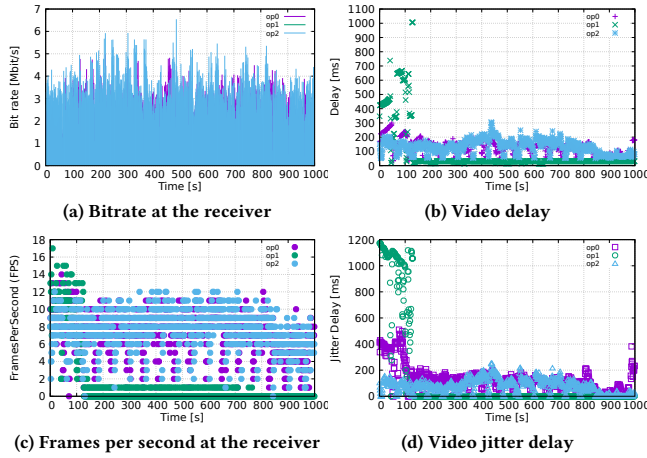
(d) Video jitter delay

Figure 8: MONROE-WebRTC performance figures (time series) for an experiment with streamers located aboard trains moving in the region of Oslo, Norway, using the three different MBB operators (one stream per train, a different operator per stream).

delay and jitter, and later suffers an almost complete connection drop (almost no delay reports were received, and only sporadic frames were reproduced). This example leads to the conclusion that current MBB networks are not ready yet to fully support WebRTC and similar video services everywhere on the move.

5 CONCLUSIONS

In this paper, we have described the MONROE platform and two containers developed to study the performance of video streaming in operational MBB networks. The MONROE-AStream and Web-Streamer containers made it possible to evaluate video streaming performance over several operator networks in multiple countries, using DASH (with different adaptation algorithms) and RTP. Our open dataset comprises not only video-related quality metrics, but also rich metadata coming from the MONROE platform that are useful to interpret the *context* of the measurements in terms of MBB parameters and allow for recreating certain scenarios.

Using the MONROE containers, we investigated the effect of different protocols and algorithms on video streaming performance. First, we evaluated DASH with different adaptation algorithms, and

second, we evaluated RTP by using WebRTC in a server-to-client video streaming mode. The measurements so far unveiled that, while mobile operators still have to strive for offering consistent quality guarantees, especially for users on the move, both DASH and RTP can provide a relatively stable performance over mobile networks when the users are stationary. We have released our complete software toolset for experiment design and data analysis—and in particular the Docker containers implementing DASH [6] and WebRTC [18]—as open source. Our rich open dataset consists of traces from video streams, performance indicators captured during the streaming sessions, and experiment metadata. Being open source, our work can be extended and reused with minor effort.

ACKNOWLEDGMENTS

REFERENCES

[1] 2016. *Technavio Global Video Services Market 2016–2020.* Technical Report.
[2] 2017. *Cisco Visual Networking Index: Forecast and Methodology, 2016–2021.* Technical Report.
[3] Ö. Alay, A. Lutu, M. Peón-Quirós, V. Mancuso, T. Hirsch, K. Evensen, A. Hansen, S. Alfredsson, J. Karlsson, A. Brunström, A. Safari Khatouni, M. Mellia, and M. Ajmone Marsan. 2017. Experience: An Open Platform for Experimentation with Commercial Mobile Broadband Networks. In *Proc. of ACM Mobicom*.
[4] AStream. 2018. (2018). https://github.com/pari685/AStream(commit:aadcfe63)
[5] ITEC Big Buck Bunny. 2018. (2018). http://www-itec.uni-klu.ac.at/ftp/datasets/mmsys12/BigBuckBunny/
[6] AStream Container. 2018. (2018). https://github.com/MONROE-PROJECT/Experiments/tree/master/experiments/astream
[7] WebStreamer Container. 2018. (2018). https://hub.docker.com/r/07777/rt/
[8] MONROE Video Dataset. 2018. (2018). https://doi.org/10.5281/zenodo.1188411
[9] FCC. 2013. *2013 Measuring Broadband America February Report.* Technical Report.
[10] S. Hu, L. Sun, C. Gui, E. Jammeh, and I. H. Mkwawa. 2014. Content-aware adaptation scheme for QoE optimized dash applications. In *2014 IEEE Global Communications Conference*. 1336–1341. https://doi.org/10.1109/GLOCOM.2014.7036993
[11] J. Huang, F. Qian, Y. Guo, Y. Zhou, Q. Xu, Z.-M. Mao, S. Sen, and O. Spatscheck. 2013. An In-depth Study of LTE: Effect of Network Protocol and Application Behavior on Performance. In *Proc. of SIGCOMM*.
[12] T. Huang, R. Johari, N. McKeown, M. Trunnell, and M. Watson. 2014. A Buffer-Based Approach to Rate Adaptation: Evidence from a Large Video Streaming Service. *SIGCOMM* (2014).
[13] P. Juluri, V. Tamarapalli, and Deep Medhi. 2015. SARA: Segment Aware Rate Adaptation Algorithm for Dynamic Adaptive Streaming over HTTP. *ICCW* (2015).
[14] MONROE manuals and publicly released code. 2018. (2018). https://github.com/MONROE-PROJECT/
[15] C. Midoglu, K. Kousias, C. Griwodz, and Ö. Alay. 2017. Evaluation of DASH Rate Adaptation Algorithms in Operational Mobile Networks (poster). *TMA PhD School* (2017).
[16] M. Moulay and V. Mancuso. 2018. Experimental Performance Evaluation of WebRTC Video Services over Mobile Networks. *INFOCOM CNERT* (2018).
[17] M. Peón Quirós, V. Mancuso, V. Comite, A. Lutu, Ö. Alay, S. Alfredsson, J. Karlsson, A. Brunstrom, M. Mellia, A. Safari Khatouni, and T. Hirsch. 2017. Results from Running an Experiment as a Service Platform for Mobile Networks. *MobiCom WiNTECH* (2017).
[18] WebStreamer repo. 2018. (2018). https://github.com/MONROE-PROJECT/Experiments/tree/master/experiments/webstreamer
[19] M. Z. Shafiq, L. Ji, A. X. Liu, J. Pang, S. Venkataraman, and J. Wang. 2013. A First Look at Cellular Network Performance during Crowded Events. In *Proc. of SIGMETRICS*.
[20] N. Vallina-Rodriguez. 2017. Illuminating the Third Party Mobile Ecosystem with the Lumen Privacy Monitor. In *FTC PrivacyCon 2017*.
[21] C. Wang, D. Bhat, A. Rizk, and M. Zink. 2017. Design and Analysis of QoE-Aware Quality Adaptation for DASH: A Spectrum-Based Approach. *ACM Trans. Multimedia Comput. Commun. Appl.* 13, 3s, Article 45 (July 2017), 24 pages. https://doi.org/10.1145/3092839

[22] A. Zeidan, A. Lehmann, and U. Trick. 2014. WebRTC enabled multimedia. In *in proceedings of World Telecommunications Congress 2014*.