

Distributed DASH Dataset

Stefan Lederer, Christopher
Mueller, and Christian
Timmerer

Alpen-Adria-Universität Klagenfurt
Universitätsstraße 65-67
9020 Klagenfurt am Wörthersee, Austria
+43 (0) 463 2700 3600
{*firstname.lastname*}@itec.aau.at

Cyril Concolato and
Jean Le Feuvre

Telecom ParisTech
46, Rue Barrault
75013 PARIS
+33 1 45 81 79 91
{*firstname.lastname*}@telecom-paristech.fr

Karel Fliegel

Czech Technical University in Prague
Faculty of Electrical Engineering
Technicka 2
166 27 Prague 6, Czech Republic
+420 224 352 248
fliegek@fel.cvut.cz

ABSTRACT

The delivery of multimedia content over HTTP and on top of existing Internet infrastructures is becoming the preferred method within heterogeneous environment. The basic design principle is having an intelligent client which selects given and applicable media representations by issuing HTTP requests for individual segments based on the users' context and current conditions. Typically, this client behavior differs between implementations of the same kind and for the objective evaluations thereof appropriate datasets are needed. This paper presents a distributed dataset for the recently published MPEG-DASH standard which is mirrored at different sites across Europe, namely Klagenfurt, Paris, and Prague. A client implementation may choose to request segments from these sites and dynamically switch to a different location, e.g., in case the one currently used causes any issues. Thus, this distributed DASH dataset can be used for real-world evaluations enabling the simulation of switching between different content delivery networks.

Categories and Subject Descriptors

H.2.4 [Systems]: Multimedia databases

General Terms

Measurement, Experimentation, Standardization, Verification.

Keywords

Dataset, Dynamic Adaptive Streaming over HTTP, DASH.

1. INTRODUCTION

The current Internet is predominated by real-time entertainment and forecasts predict it will be around 70% in 2017 in both wired and wireless environments [1]. Most of its traffic is dedicated to video referring to visual, audio, and associated text (e.g., subtitles, closed caption) information delivered over the Hypertext Transfer Protocol (HTTP).

In recent years the delivery of video over HTTP has gained momentum and various proprietary formats (e.g., Adobe HTTP Dynamic Streaming, Apple HTTP Live Streaming, Microsoft

Smooth Streaming) evolved to an international standard referred to as MPEG Dynamic Adaptive Streaming over HTTP (DASH) [2]. In general, the client receives a manifest describing the relationship of media representations (e.g., different bitrates, resolutions, codecs, languages) and media timeline for which HTTP requests are issued to fetch individual segments according to the conditions measured at the client. The client implementation itself is non-normative and left open for industry competition. In particular, part one of MPEG-DASH specifies formats for an XML-based Media Presentation Description (MPD) representing the manifest and segment formats based on the ISO base media file format (ISO/BMFF) and MPEG-2 Transport Stream (M2TS) [3].

In practice, there is a need for public datasets that enable an objective comparison of evaluation results across different client implementations within both academic and industrial environments. Multiple efforts are undertaken for creating this kind of datasets such as the QUALINET database [4] and the MMSys dataset track [5]. A first dataset for MPEG-DASH is provided in [6] comprising various full-length videos in a variety of configurations in terms of genre, resolution, bitrate, segmentation, and segment length. It has been developed in the course of the MPEG-DASH standardization and is constantly updated to be compliant to the final version of the standard including its amendments and corrigenda. The latest version is available under http://www-itec.uni-klu.ac.at/dash/?page_id=207. Additionally, the GPAC project at Telecom ParisTech offers synthetic DASH sequences for the purpose of conformance testing which are available under <http://gpac.wp.mines-telecom.fr/2012/02/23/dash-sequences/>. In this context, MPEG is also working towards finalizing part two of DASH referred to as conformance and reference software [7] providing software and test vectors in order to test whether implementations are compliant with the standard. Finally, the recently established DASH Industry Forum (DASH-IF) [8] is also working towards test vectors similar to MPEG and going beyond by including also guidelines for specific codecs.

In this paper we provide a dataset which enables dynamic, adaptive HTTP streaming according to MPEG-DASH with multiple *BaseURLs*. That is, a full-length DASH video from [4] is replicated on three sites within Europe (i.e., Klagenfurt, Paris, and Prague) which is indicated within the MPD by multiple *BaseURLs* allowing for a real-world evaluation of DASH clients that perform bitstream switching between multiple sites. Please note that this could be useful as a simulation of switching between multiple Content Distribution Networks (CDNs). In particular, it may be

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, or 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.

used to evaluate different MPEG-DASH client implementations under real-world conditions to allow an objective comparison thereof. Additionally, we provide a mechanism to mirror the DASH content to further sites while maintaining consistency of the MPD concerning the usage of the *BaseURL*.

The remainder of the paper is organized as follows. Section 2 describes the usage of multiple BaseURLs within MPEG-DASH whereas Section 3 describes the actual dataset, its distribution mechanism, and the mirror sites. Finally, research challenges related to this dataset are highlighted in Section 4 and Section 5 concludes the paper.

2. DASH AND MULTIPLE BASEURLS

Part one of the MPEG-DASH standard [3] introduces the *BaseURL* element as part of the *Media Presentation Description* (MPD) syntax. The *BaseURL* element specifies, in its text content, a URL indicating a location that can be used to request the different segments needed for the presentation: media segments, initialization segments, index segments, or bitstream switching segments. This is an optional element which can be present multiple times at multiple levels in the XML hierarchy of the MPD.

When multiple *BaseURL* elements are specified at a given level in the MPD hierarchy, as shown in Listing 1, this means that the different URLs are alternatives that can be used by the client. Conformance to the MPEG-DASH standard requires that the same media segments be available on all servers. Please note that no specific order, priority, or preference between the URL alternatives is indicated in the MPD. The alternative server locations may be used as fallback when one server becomes unavailable. A client may also decide to switch the location based on network bandwidth or round trip time (RTT).

```
<MPD xmlns="urn:mpeg:dash:schema:mpd:2011" ...>
  <Period ...>
    <AdaptationSet ...>
      <Representation ...>
        <BaseURL>http://server1.example.org/path/</BaseURL>
        <BaseURL>http://server2.example.net/path/</BaseURL>
        <SegmentList ...> ...
          <SegmentURL media="path/segment1.m4s"/>
          <SegmentURL media="path/segment2.m4s"/>
          <SegmentURL media="path/segment3.m4s"/>
          ...
        </SegmentList>
      </Representation>
    </AdaptationSet>
  </Period>
</MPD>
```

Listing 1. Example of multiple BaseURL elements at the same level in the MPD hierarchy.

When multiple *BaseURL* elements are used at different levels in the MPD hierarchy, as depicted in Listing 2, the URL to be used to retrieve the segments is build following the resolution algorithm specified in RFC3986 [9]. The resolution algorithm describes how to merge the URLs from different levels in the hierarchy. In particular, relative URLs, i.e., those not starting with a protocol scheme such as “http://” or “https://”, are resolved with the *BaseURL* on the level above, which may be relative also. In the special case where the top level resolved URL is also relative, it is finally resolved with the URL of the MPD, thereby indicating

that the media segments are located at the same server as the MPD.

```
<MPD xmlns="urn:mpeg:dash:schema:mpd:2011" ...>
  <BaseURL>http://server1.example.org/</BaseURL>
  <Period ...>
    <AdaptationSet ...>
      <Representation ...>
        <BaseURL>path_to_representation/</BaseURL>
        <SegmentList ...> ...
          <SegmentURL media="path/segment1.m4s"/>
          <SegmentURL media="path/segment2.m4s"/>
          <SegmentURL media="path/segment3.m4s"/>
          ...
        </SegmentList>
      </Representation>
    </AdaptationSet>
  </Period>
</MPD>
```

Listing 2. Example of multiple BaseURL elements at different levels in the MPD hierarchy.

The *BaseURL* element may have two optional attributes called *byteRange* and *serviceLocation*. The *serviceLocation* attribute is a string indicating that the URL specified in the *BaseURL* element belongs to the same CDN as other *BaseURL* elements with the same value of *serviceLocation*. The *byteRange* attribute specifies a template to construct HTTP URL based on the byte range of the segments.

3. DISTRIBUTED DASH DATASET

3.1 Dataset Description

The RedBull Playstreet sequence as used in [4] has been adopted as a basis for this dataset. The sequence has a length of 1 hour, 37 minutes and 28 seconds containing both audio and video.

The DASH content is provided with different segment lengths of 2, 4, 6, 10, and 15 seconds. This becomes necessary as the performance of DASH-based streaming is influenced by network characteristics such as the round trip time (RTT) as shown in previous work [4] as well as HTTP server configurations like the usage of persistent HTTP connections or pipelining [10]. If shorter segment lengths are used, the influence of the network delay gets bigger as when the client sends HTTP requests for segments and the downlink is not utilized in the meantime. The above mentioned optimization possibilities aim to reduce this time where the downlink is not fully utilized, like it is done in [10] using one persistent TCP connection over the whole streaming session and pipelined segment requests, and thus gain a significantly better streaming performance.

The different bitrates and resolutions of the dataset are shown in Table 1 where for each segment length the content was encoded at 17 different video representations, ranging from 100kbps at 320x240 up to 6 Mbps at 1920x1080. The lower bitrates and resolutions would be suitable for mobile devices like smartphones with limited network connectivity whereas the higher bitrates and resolutions could be used for TV sets, set-top boxes, or personal computers. Additionally, four different audio representations are provided with two channels at 64, 96, 128, and 165 kbps using a 48 kHz sampling rate. The audio and video representations are offered separately enabling the client to choose the appropriate audio and video bitrate independently from each other, e.g., based on the users' context like device and display type, user preferences, etc.

Table 1. Bitrates and resolutions for video and audio.

#	Video (AVC)	Audio (AAC)
1	100 kbit/s, 320x240	64 kbit/s, 48 kHz, 2 channel
2	150 kbit/s, 320x240	96 kbit/s, 48 kHz, 2 channel
3	200 kbit/s, 480x360	128 kbit/s, 48 kHz, 2 channel
4	250 kbit/s, 480x360	165 kbit/s, 48 kHz, 2 channel
5	300 kbit/s, 480x360	
6	400 kbit/s, 480x360	
7	500 kbit/s, 854x480	
8	700 kbit/s, 854x480	
9	900 kbit/s, 854x480	
10	1,2 Mbit/s, 854x480	
11	1,5 Mbit/s, 1280x720	
12	2,0 Mbit/s, 1280x720	
13	2,5 Mbit/s, 1280x720	
14	3,0 Mbit/s, 1920x1080	
15	4,0 Mbit/s, 1920x1080	
16	5,0 Mbit/s, 1920x1080	
17	6,0 Mbit/s, 1920x1080	

The dataset was produced using tools from the GPAC project [11], in particular with MP4Box including the DASH segmenter and MPD writer. The video representations have been encoded using x264 [12] and the audio representations using the AAC encoder provided by FFmpeg [13]. The segment format is based on the ISO/BMFF and the total size of the dataset is approximately 85 GB.

The actual MPDs are offered in two versions:

- The first version adopts the *segment approach* where each media segment is indexed separately within the MPD. For this version the profile `urn:mpeg:dash:profile:isoff-main:2011` of the DASH standard was used.
- As the first version may result in increased MPD sizes, the second version is leveraging *URLTemplates* which reduces the size of the MPD significantly. The segments are identified by a naming scheme and not referenced separately. The profile used for the *URLTemplate*-based MPDs is `urn:mpeg:dash:profile:isoff-live:2011` which has been selected in version 1.5 of the HbbTV specification [14]. Hence, it also offers the possibility to do simulations based on live content.

The MPDs have been validated with the DASH MPD validator (http://www-itec.uni-klu.ac.at/dash/?page_id=605). Playback has been tested with various DASH implementations, including GPAC, DASH VLC Plugin, libdash, and DASH-JS.

3.2 Main Repository and Distribution

The *main repository* of the dataset is located at the Alpen-Adria-Universität (AAU) Klagenfurt and is available under the following the URL:

<http://www-itec.aau.at/ftp/datasets/mmsys13/>

It is mirrored to the other sites participating in the dataset and described below. Each *BaseURL* in the MPDs points to a location comprising a mirror site of the dataset. Additionally, one *BaseURL* provides the relative URL to the local server indicated by “/”. Hence, a DASH client has the possibility to choose one of the preferred sites or dynamically switches between the available sites based on the actual bandwidth measurements during the streaming session.

The dataset is open for new participants according to the details provided at http://www-itec.uni-klu.ac.at/dash/?page_id=958. This Web site also shows the status of all participating sites:

- **registered** for new participants whose mirror needs to be verified;
- **active** for verified mirrors up and running;
- **inactive** for sites which are unavailable, have incomplete replicas, or are not up-to-date.

The progress of adding new participants who want to join the dataset is as follows:

1. The dataset has to be mirrored to the new location based on the URL of the main repository given above or using the FTP URL: `ftp://ftp-itec.aau.at/pub/datasets/mmsys13/`. Furthermore, a mechanism for keeping the mirror synchronized with the main repository shall be applied to guarantee up-to-date MPDs at the new site.
2. The new mirror location has to be registered at the DASH Web site at AAU, requiring information about the institution, the contact person and its email address, and the *BaseURL* of the new dataset replica.
3. After the successful registration of the new site its status is set to *registered*. Please note that its *BaseURL* is not yet included within the MPDs of the dataset until it has been verified in terms of availability, correctness of the *BaseURL*, completeness of the dataset replica, etc. If everything is correct the status is changed to *active* and the site is added to the *BaseURLs* of the dataset MPDs within the main repository at AAU. As each mirror leverages a replication mechanism, these MPDs are also updated on each site of the dataset after a certain time span (ideally one day).
4. All mirrors of the dataset are checked periodically (using Nagios: <http://www.nagios.org/>) for their availability and completeness. Unavailable or incomplete mirrors will be set to inactive in the *BaseURL* database, which has the consequence, that they are not included in the *BaseURLs* of the dataset MPDs anymore. In this case the site administrator will be contacted to fix the problems.

In addition to the MPDs of the dataset, we provide a MPD generation service under the following URL:

<http://www-itec.aau.at/dash/ddash/mpdGenerator.php>

It offers the possibilities to configure the requested MPD via HTTP GET parameters *segmentlength* and *type*:

- The parameter *segmentlength* is used to request the MPD for a given segment length with the following scope `segmentlength={2, 4, 6, 10, 15}`. Each number specifies the segment length in seconds.
- The parameter *type* is used to define the structure of the MPD with the following scope `type={full, URLTemplate}`. The former is used for MPDs referencing each segment separately and the latter for MPDs with the URL template mechanism.

In both cases when the parameter values are outside its scope, a HTTP 404 error is returned.

The following request is an example requesting an MPD describing the dataset content with two seconds segment length and referencing each media segment separately: `http://www-itec.aau.at/dash/ddash/mpdGenerator.php?segmentlength=2&type=full`.

3.3 Additional Repositories

The main repository is currently mirrored at two further locations within Europe, namely Paris and Prague, and described in the following.

The first mirror location of the dataset proposed in this paper is at Telecom ParisTech, Paris, France. The content is available under the following URL:

<http://download.tsi.telecom-paristech.fr/gpac/dataset/dash/mmsys13/>

It is mirrored daily at midnight from the main repository. The mirror has been done using *wget* with the following command line:

```
wget -m -nH --cut-dirs=3 ftp://ftp-itec.aau.at/pub/datasets/mmsys13/
```

The second mirror location of the dataset proposed in this paper is at Czech Technical University in Prague, Czech Republic. The content is available under the following URL:

<http://dbq.multimediatech.cz/mmsys13/>

It is mirrored daily at midnight from the main repository. The dataset is part of the QUALINET DATABASES [4]. One of the goals within the interests of the COST Action IC1003 QUALINET and its Working Group 4 (Databases and Validation) is to create a database of proper multimedia content and take the steps to make it accessible to all researchers. This database is available to the scientific community through a simple Web portal. Registered users (QUALINET members and approved users from the scientific community) have access to the list of multimedia databases and based on their user rights they are allowed to browse information about the particular database and eventually download the actual multimedia content. Additionally, the Web site provides a wiki page which informs about the implementation process of the QUALINET DATABASES platform. At the time of writing this paper, there are about 110 datasets listed in the QUALINET DATABASES, from which about half was created and is owned by QUALINET partners. The described databases can be divided into three main categories: (1) annotated multimedia quality databases, (2) eyetracking databases, and (3) other databases. The databases with the direct utilization in Quality of Experience (QoE) testing and benchmarking are mainly from (1) annotated multimedia quality databases. The further classification of the datasets is based on the actual content type: (1.1) annotated image quality databases and (1.2) annotated video quality databases. These databases of images or videos are annotated with the results from subjective tests. There is also special multimedia content available such as 3D images and videos or results from visual attention experiments. Detailed description of the datasets available within QUALINET DATABASES is available at the online platform [4].

4. RESEARCH CHALLENGES

In this section we highlight research challenges for (client) implementations when faced with multiple *BaseURLs* present within an MPD. Obviously, the first challenge when retrieving an MPD with multiple *BaseURLs* is determining with which *BaseURL* to start a DASH session. As the *BaseURL* does not have any metrics associated (except the service location which could be seen as a kind of grouping mechanism) it is up to client implementation to decide the location of the first segments to be downloaded. In any case, finding out the best *BaseURL* to use may influence the (start-up) delay.

The second challenge is due to possible bandwidth fluctuations during a DASH session, which raises the issue whether the client shall switch to another *BaseURL* (probably from a different service location, if available) or stay within the same *BaseURL* and simply select another representation. In other words, it is not always obvious whether bandwidth fluctuations occur temporarily which suggests switching to another representation or indicate a permanent issue which suggests switching to another *BaseURL*. If the decision is to switch to another *BaseURL*, the client needs to decide to which *BaseURL* it should switch. This decision is actually similar to the first challenge but at a different time scale.

In general we argue that estimating the bandwidth for segments located at multiple *BaseURLs* in parallel while providing a low start-up delay and smooth streaming without stalls or re-buffering is a major critical issue for DASH implementations. Additionally, in case the delivery network comprises proxy caches (which is most likely the case) where multiple clients compete for a bottleneck, accurate bandwidth estimation becomes even more complex, like shown in [15].

A third challenge is related to live streaming and covers a plethora of issues ranging from live distribution of MPDs (and keeping them up-to-date) and segments, selection of the nearest/best *BaseURL* for the actual streaming, and handling of server errors probably followed by choosing another *BaseURL*. This is specifically challenging as for live services the delay and time scale is usually much shorter than for on demand services where the content is distributed much more in advance. In this context, simulations of such scenarios can be conducted based on the MPDs and different locations of the dataset.

5. CONCLUSIONS

In this paper we have presented a distributed DASH dataset comprising multiple *BaseURLs*. It is complementary to the existing DASH dataset [4] and provides a unique property of being distributed across three sites within Europe which is indicated by multiple *BaseURL* elements within the MPD. Hence, the dataset is applicable for performing real-world experiments with each site simulating a different CDN location and bitstream switching across multiple CDNs. Additionally, we provide a mechanism for further distributing the dataset and a centralized tracking service as an entry point for the entire distributed DASH dataset. Future work will focus on the promotion of this dataset within both academic and industrial environments as well as on addressing the research challenges identified in this paper.

6. ACKNOWLEDGMENTS

This work was supported in part by the EC in the context of the ALICANTE (FP7-ICT-248652) and SocialSensor (FP7-ICT-287975) projects and partly performed in the Lakeside Labs research cluster at AAU. Special thanks to the Red Bull Media House for providing us the Red Bull Playstreets video. They own the rights of the content but the usage for scientific purposes is permitted. This work was also supported in part by the French-funded project AUSTRAL (DGCIS FUI13). This work was partially supported by the COST IC1003 QUALINET, by the Czech-funded project COST CZ LD12018 MOVERIQ and by the grant of the Czech Science Foundation No. P102/10/1320.

7. REFERENCES

- [1] Sandvine 2012, *Global Internet Phenomena Report 1H 2012*, Sandvine Intelligent Broadband Networks, <http://www.sandvine.com/downloads/documents/Phenomena>

- _1H_2012/Sandvine_Global_Internet_Phenomena_Report_1H_2012.pdf (last access: Nov. 2012).
- [2] Sodagar, I. 2011. The MPEG-DASH Standard for Multimedia Streaming Over the Internet. *IEEE MultiMedia* 18, 4 (Oct. 2011), 62-67. DOI=<http://dx.doi.org/10.1109/MMUL.2011.71>
- [3] ISO/IEC 23009-1:2012, *Information technology — Dynamic adaptive streaming over HTTP (DASH) — Part 1: Media presentation description and segment formats*. http://standards.iso.org/ittf/PubliclyAvailableStandards/c057623_ISO_IEC_23009-1_2012.zip (last access: Nov. 2012)
- [4] QUALINET DB, <http://dbq.multimediatech.cz/> (last access: Nov. 2012).
- [5] MMSys Dataset Track, <http://web.cs.wpi.edu/~claypool/mmsys-dataset/> (last access: Nov. 2012).
- [6] Lederer, S., Mueller, C., and Timmerer, C. 2012. Dynamic adaptive streaming over HTTP dataset. In *Proceedings of the 3rd Multimedia Systems Conference (MMSys '12)*. ACM, New York, NY, USA, 89-94. DOI=<http://doi.acm.org/10.1145/2155555.2155570>
- [7] ISO/IEC DIS 23009-2, *Information technology — Dynamic adaptive streaming over HTTP (DASH) — Part 2: Conformance and reference software*.
- [8] DASH Industry Forum, <http://www.dashif.org/> (last access: Nov. 2012).
- [9] Berners-Lee, T., Fielding, R., and Masinter, L. 2005. Uniform Resource Identifier (URI): Generic Syntax. *RFC3986*, <http://www.ietf.org/rfc/rfc3986.txt> (last access: Nov. 2012).
- [10] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and Berners-Lee, T. 1999. Hypertext Transfer Protocol -- HTTP/1.1. *RFC2616*, <http://www.w3.org/Protocols/rfc2616/rfc2616.html> (last access: Nov. 2012).
- [11] Le Feuvre, J., Concolato, C., Dufourd, J.-C., Bouqueau, R., and Moissinac, J.-C. 2011. Experimenting with multimedia advances using GPAC. In *Proceedings of the 19th ACM international conference on Multimedia (MM '11)*. ACM, New York, NY, USA, 715-718. DOI=<http://doi.acm.org/10.1145/2072298.2072427>
- [12] X264, <http://www.videolan.org/developers/x264.html>, (last access: Nov. 2012).
- [13] FFMPEG, <http://www.ffmpeg.org/>, (last access: Nov. 2012).
- [14] HbbTV. 2012. HbbTV® Specification Version 1.5. http://www.hbbtv.org/pages/about_hbbtv/HbbTV-specification-1-5.pdf (last access: Nov. 2012).
- [15] Mueller, C., Lederer, S. and Timmerer, C. 2012. A Proxy Effect Analysis and Fair Adaptation Algorithm for Multiple Competing Dynamic Adaptive Streaming over HTTP Clients. In *Proceedings of the IEEE Conference on Visual Communications and Image Processing Conference (VCIP 2012)*. IEEE, San Diego, CA, USA, pp. 6, 2012.