

Dynamic Adaptive Streaming over HTTP

FROM CONTENT CREATION TO CONSUMPTION

ACM MULTIMEDIA
Tutorial, Oct/Nov, 2012
Nara, Japan
www.acmmm12.org

Christian Timmerer (Alpen-Adria-Universität Klagenfurt, Austria) | Carsten Griwodz (Simula Research Laboratory, Norway)

Cover Sheet

Title: Dynamic Adaptive Streaming over HTTP – From Content Creation to Consumption

Target length of the tutorial: half day


The intended audience: intermediate, graduate/PhD students and researchers interested in the delivery of multimedia content using the recently ratified MPEG-DASH standard.

Complete contact information for the contact person and other presenters:

- Christian Timmerer, Universitätsstraße 65-67, Klagenfurt, 9020, Austria, Tel: +43 463 2700 3621, Fax: +43 463 2700 99 3621; E-mail: christian.timmerer@itec.aau.at, Web: <http://research.timmerer.com> Blog: <http://blog.timmerer.com>, Twitter: <http://twitter.com/timse7>
- Carsten Griwodz, Martin Linges v 17, Fornebu, Norway, Tel: +47 45258004, Fax: +47 67 82 82 01; Email: griff@simula.no, Web: <http://simula.no/people/griff>

Biography of Presenters

Christian Timmerer is an assistant professor at the Institute of Information Technology (ITEC), Multimedia Communication Group (MMC), Alpen-Adria-Universität Klagenfurt, Austria. His research interests include the transport of multimedia content, multimedia adaptation in constrained and streaming environments, distributed multimedia adaptation, and Quality of Service/Quality of Experience. He was the general chair of WIAMIS'08, AVSTP2P'10 (co-located with ACM MM'10), WoMAN'11 (co-located with ICME'11), and TPC chair of QoMEX'12. He has participated in several EC-funded projects, notably DANAE, ENTHRONE, P2P-Next, ALICANTE, and SocialSensor. He is an Associate Editor for IEEE Computer Science Computing Now, Area Editor for Elsevier Signal Processing: Image Communication, Review Board Member of IEEE MMTC, editor of ACM SIGMM Records, and member of ACM SIGMM Open Source Software Committee. He also participated in ISO/MPEG work for several years, notably in the area of MPEG-21, MPEG-M, MPEG-V, and DASH. He received his PhD in 2006 from the Klagenfurt University. Publications and MPEG contributions can be found under <http://research.timmerer.com>, follow him on <http://www.twitter.com/timse7>, and subscribe to his blog <http://blog.timmerer.com>. Full bio can be found at <http://www-itec.uni-klu.ac.at/~timse/cv/>.



Carsten Griwodz is head of the Media department of government-owned research company Simula Research Laboratory, and professor of Computer Science at the University of Oslo. He received his Dipl.-Inf. degree from Paderborn University in 1993 and Dr.-Ing. degree from Technische Universität Darmstadt in 2000. He worked for IBM from 1993–98 and participated in the standardization of MHEG. His research is concerned with streaming media, ranging from scalable distribution architectures through operating system and protocol support to subjective visual quality assessment. He was co-chair of ACM NOSSDAV 2008, ACM/IEEE NetGames 2011, SPIE/ACM MM CN 2006 and 2007, Track chair of ACM MM 2008, TPC chair of ACM MMSys 2012 and is general chair of MMSys 2013. He is Associate Editor of ACM TOMCCAP and Editor-in-Chief of the newsletter ACM SIGMM Records. The Media group publishes news at <http://mpg.ndlab.net>. His publications can be found at <http://simula.no/people/griff/bibliography>.



Dynamic Adaptive Streaming over HTTP – From Content Creation to Consumption

(Extended Abstract)

Standards developing organizations (SDOs) such as MPEG have developed various technologies for multimedia transport, e.g., MPEG-2 TS (Transport Stream) and the MP4 file format. These technologies have been widely accepted and are heavily used by various industries and applications, such as digital broadcasting, audio and video transport over the Internet and streaming to mobile phones, etc. At the same time, many other SDOs such as the IETF, IEEE, and 3GPP have been providing various protocols to deliver multimedia content packetized or packaged by such MPEG transport technologies.

Broadcasting services and mobile services are converging, and it is expected that this convergence trend will continue with other services. Additionally, new emerging multimedia services are being introduced. These developments in the multimedia arena mean that various content and services will be delivered over different networks, and the users expect to consume these services using those networks, depending on the availability and reach of the network at the time of consumption. To deploy efficient solutions for the transport of modern media in an interoperable and universal fashion, especially given the recent increased demand in the heterogeneous network environment, there is urgency for an international multimedia transport standard.

In recent years, the Internet has become an important channel for delivery of multimedia. The HTTP protocol is widely used on the Internet. Recently, it has also become a primary protocol for the delivery of multimedia content, and a number of proprietary solutions are available. Several of these solutions implement adaptive streaming over HTTP, which allows clients that receive a video to observe their own network throughput and use this information to choose from several levels of compression that are made available by servers. This quality switching can be done at predefined points in time of the video's playback schedule.

The advantage of this approach is that it relies on widely deployed infrastructure and a pragmatic division of concerns. However, dynamic adaptive streaming is far less understood than classical streaming approaches, and much research is still required to identify how components interact with each other, how alternative compression techniques can be used, how it can be used in the context of applications, and how overall performance can be optimized.

The aim of this tutorial is to provide an overview of the recently ratified MPEG-DASH standard, how to create content to be delivered using DASH, its consumption, and the evaluation thereof with respect to competing industry solutions. One of the essential differences between DASH and earlier streaming solutions is that the client is in charge of adapting its demands to the bandwidth share that is available for serving it. This client behavior is not standardized, and differentiates players. We are going to present how current industry standard solutions, including Smooth Streaming, Netflix's variant, HLS and Adobe perform; we explain how this performance comes to pass and where each of the solutions is applicable. It may not be obvious, but widely different behavior is desirable for on-demand streaming to wireless devices, on-demand streaming to home cinemas, joining continuous live streams, or waiting for a live event. We present the differences and the best known adaptation decisions. Obviously, the strategies matter only if network performance varies. We will therefore also explain how different bottlenecks in the network can be approached.

DASH is an excellent format for non-linear consumption; it excels through its simplicity. We will discuss how the dynamic creation of DASH manifest files (and equivalent industry standard approaches) can be used for summarization or editing. We'll also discuss some known shortcomings in live streaming.

Viewers of a video stream experience of good quality of experience when they watch video in high quality immediately after choosing to do so, without stalls and with good visual quality. Adaptive can provides this (most of the time) without requiring the user to make choices, which makes the service provided more attractive. However, there is a problem in increasing and decreasing qualities as networks' bandwidth presents itself. The quality change itself is disruptive to user experience; encodings can be changed along several parameters (framerate, picture resolution, sharpness), which have different properties; enhancing quality at the price of an increased startup delay improves user experience in some situations. We discuss the variety of options and present state-of-the-art knowledge for making these choices.

Course Outline

The tutorial can be roughly clustered into three parts. In part I we will provide an introduction to DASH, part II covers content creation, delivery, and consumption, and, finally, part III deals with the evaluation of existing (open source) MPEG-DASH implementations compared to state-of-art deployed industry solutions.

Introduction to Dynamic Adaptive Streaming over HTTP (DASH)

- Motivation: why HTTP streaming?
- Background on MPEG-DASH standardization.
- MPEG-DASH, the enabler for interoperable streaming over HTTP.
- Media Presentation Description (MPD), segment formats (ISOBMFF, M2TS, [WebM]), profiles, deployments (DASH Promoters Group).
- Applications requirements: video-on-demand, (quasi-) live streaming, video editing, summarization, video search.

Content Creation, Delivery, and Consumption

- How to create content using DASH: coding format (AVC vs. SVC), bitrate, resolution, framerate, segment length, quality.
- How to delivery content using DASH: content delivery network and proxy cache issues.
- How to consume content using DASH: client issues featuring adaptation logic, bandwidth estimation, segment request scheduling.
- Open source implementation efforts: VLC, GPAC, libdash, DASHEncoder, dataset, DASH-JS.

Evaluation

- Industry solutions (Microsoft, Apple, Netflix, Adobe) vs. MPEG-DASH implementations.
- Quality of Service (QoS) / Quality of Experience (QoE).