Adaptive Multimedia Streaming in Information-Centric Networks

Stefan Lederer, Christopher Mueller, Christian Timmerer, and Hermann Hellwagner

Abstract

ICN has received a lot of attention in recent years, and is a promising approach for the Future Internet design. As multimedia is the dominating traffic in today's and (most likely) the Future Internet, it is important to consider this type of data transmission in the context of ICN. In particular, the adaptive streaming of multimedia content is a promising approach for usage within ICN, as the client has full control over the streaming session and has the possibility to adapt the multimedia stream to its context (e.g. network conditions, device capabilities), which is compatible with the paradigms adopted by ICN. In this article we investigate the implementation of adaptive multimedia streaming within networks adopting the ICN approach. In particular, we present our approach based on the recently ratified ISO/IEC MPEG standard Dynamic Adaptive Streaming over HTTP and the ICN representative Content-Centric Networking, including baseline evaluations and open research challenges.

Internet traffic is dominated by multimedia data and is producing 62 percent of the total Internet traffic in North Americas' fixed access networks and will continue to grow in the next years [1]. This trend presents problems for content provides as well as Internet service providers (ISPs), as the used protocols have not been designed for multimedia delivery, which causes congested networks as well as high infrastructure needs and costs. The majority of this multimedia data is coming from popular streaming service providers like Netflix or YouTube. These service providers typically leverage existing Internet infrastructures based on the Transmission Control Protocol (TCP) and the Internet Protocol (IP). In this context the adaptive delivery of multimedia content over the cost-efficient and scalable Hypertext Transfer Protocol (HTTP) is gaining increasing momentum and has resulted in the standardization of MPEG Dynamic Adaptive Streaming over HTTP (DASH) [2]. MPEG-DASH currently receives much attention and offers a company-independent streaming standard for various environments (fixed, mobile) and devices like desktops, smartphones, tablets, and (smart) set-top boxes/televisions, which will be the major clients also in the Future Internet.

Existing networks and systems are evolutional improvements of the Internet as we know it, which is mainly based on the host-to-host connection paradigm originated in the 1970s when multimedia was not as popular as it is today. A variety of revolutionary Internet architectures have been proposed in the last decade [3], aiming to address the needs of current and future networks. Some seem to overcome current limitations, and one of these new Internet architectures is referred to as Information-Centric Networking (ICN), with representatives like Content-Centric Networking (CCN) [4]. ICN moves the focus from traditional end-to-end connections to the content and the addressing thereof using content naming schemes rather than addressing its location, that is, nodes within a network.

Given the popularity of adaptive multimedia streaming (AMS) over HTTP and the rise of ICN as a promising candidate for the Future Internet architecture, it is worthwhile to investigate whether ICN is ready for current multimedia streaming approaches and vice versa. This research is becoming increasingly important as current AMS deployments try to mimic ICN functionality through advanced content distribution networks (CDNs) and HTTP/DNS interactions without considering a deeper integration of both technologies. As ICN and AMS have several elements in common, for example a receiver-driven approach, the content being dealt with in pieces, and the usage of caching, we believe this will be a promising combination.

The purpose of this article is to present adaptive multimedia streaming in information-centric networks using two commonly acknowledged and concrete approaches, namely MPEG-DASH and CCN. In particular, we describe the architecture to enable DASH over CCN, available open source tools and test-beds, and a baseline evaluation thereof. Finally, we highlight research challenges and open issues associated with this proposal.

Information-Centric Networking

In contrast to IP, where content is explicitly exchanged among nodes, ICN is directly requesting content pieces from the network without addressing its location. When clients want to consume a specific content, they simply send out an interest for this content, which is identified by a uniform resource identifier (URI) using a routable name scheme, for example, /example.com/video/example.mp4.

ICN-approaches such as CCN and its open source implementation CCNx introduce advanced routers with caching support, which check whether the requested content is already

The authors are with Alpen-Adria-Universität Klagenfurt.

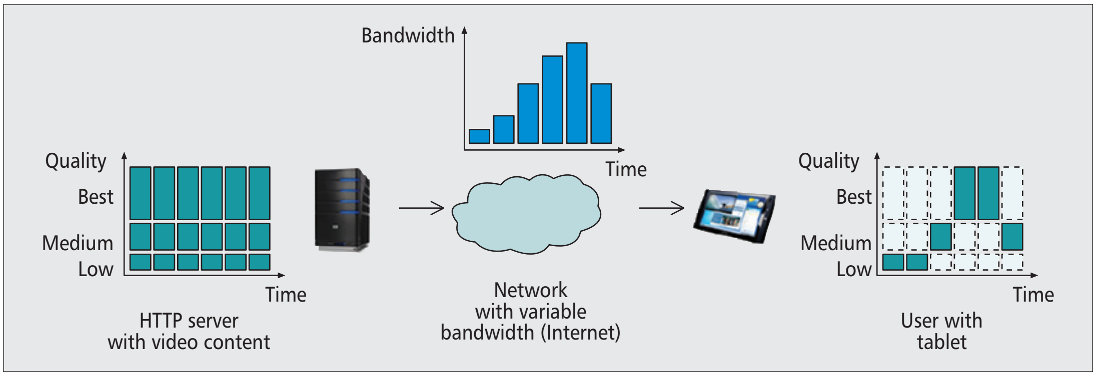

Figure 1. Adaptive multimedia streaming principles.

cached in the so-called content-store (CS) or whether there is another pending interest in the routers' pending interest table (PIT) for the same content. In case there is already a forwarded interest for a specific content, further interests for this content are not forwarded and stored in the PIT instead. When the requested content arrives at the router, all pending interests for this content are satisfied. Thus, the network provides implicit support for multicast, which is specifically interesting for live video streaming. The forwarding of an interest toward the content location is done via one or more interfaces, based on routes to the origin node(s) provided by the routers' forwarding information base (FIB). In addition to the routes in the FIB, a so called strategy layer decides on which routes and on which interfaces the interest is forwarded. As this decision is done on a per-interest level, a seamless handover between interfaces is provided, which is specifically valuable in mobile environments where devices are equipped with multiple interfaces such as 3G, 4G, and WiFi. This can also be seen as intrinsic error resilience with respect to the network. Additionally, approaches like CCN are meant to provide security und trust as an integral part of the network.

Adaptive Multimedia Streaming over HTTP

The early days of streaming multimedia over HTTP began with progressive download, where the client opens a TCP connection to a server and progressively downloads the multimedia content. As soon as enough data is available on the client, the client could eventually start with the decoding and the rendering, respectively. At this time, video streaming was mainly associated with RTP/UDP — see for example [5] for mobile environments using 3GPP Packet-Switched Streaming Service — but the usage of HTTP turned out to be an effective solution, solving problems such as passing firewalls or network address translation traversal. Additionally, it can be built upon cost-effective HTTP-based infrastructure including CDNs and proxy caches. Using progressive downloading, however, clients cannot react in case of bandwidth fluctuations which is typically followed by (re-)buffering periods also referred to as stalls. One of the first adaptive multimedia streaming over HTTP solutions employed an explicit adaptation loop where clients perform bandwidth measurements and push the information toward the server [6], which can also be found in RTP-based systems. The server analyzes these reports and modifies the progressive download session on-demand.

Today there are various industry solutions available, each following the common principle of requesting chunked multimedia data available in different qualities from conventional Web servers, instead of dedicated streaming servers. The concept of adaptive multimedia streaming over HTTP is depicted in Fig. 1. The media content will be encoded in different versions providing a variety of bitrates, resolutions, codecs, and so on. These versions are chopped into segments of a given length, which are requested individually by the client using HTTP GET requests following a pull-based approach. The content is transmitted over the Internet with varying bandwidth conditions and the client may adapt the streaming session to prevent stalls of the media playback. This is typically achieved by downloading the appropriate quality version of the segmented content. Seamless switching between these versions is typically achieved thanks to time-aligned segments that are encoded independent of each other, for example, starting with an intra-frame. Because the logic of such systems is located at the client and the stateless design of HTTP, it is possible to leverage existing HTTP infrastructures at scale.

MPEG-DASH [2] provides an open standard comprising the specification of an XML-based manifest referred to as Media Presentation Description (MPD) which describes the relationship between the segment locations using HTTP URLs with its associated characteristics (e.g. bitrate, resolution, codec) and timeline. With the information contained in the MPD, the DASH client is able to start the streaming session and adapt to bandwidth fluctuations.

Adaptive Multimedia Streaming in ICN: DASH over CCN

In this section we will describe the integration of adaptive multimedia streaming in ICN. Both approaches have several elements in common, e.g. they facilitate a receiver-driven approach using segmented data and leverage the efficient distribution and caching of the data within the network. Since those two approaches are located at different protocol layers, i.e. adaptive multimedia streaming at the application layer and ICN at the network layer, they can be combined efficiently to leverage the advantages of both.

There are two major options to combine adaptive multimedia streaming and ICN: a proxy service between HTTP and ICN as proposed in [7], and an adaptive streaming client implementing a native ICN interface. A proxy service translates conventional HTTP requests to corresponding ICN requests and returned data packets to HTTP responses. In this article we are focusing on a more integrated approach aiming at fully exploiting the potential of ICN by using a native ICN interface within the adaptive streaming client. Therefore, we make use of concrete implementations of both technologies, i.e. DASH and CCN. The former is an international standard for which comprehensive open source tools are available. The latter is an example of an ICN that we used as the basis for our work. In the following we describe our proposed DASH over CCN architecture as well as open source tools and datasets using the existing open source CCNx implementation (www.ccnx.org) as a basis.

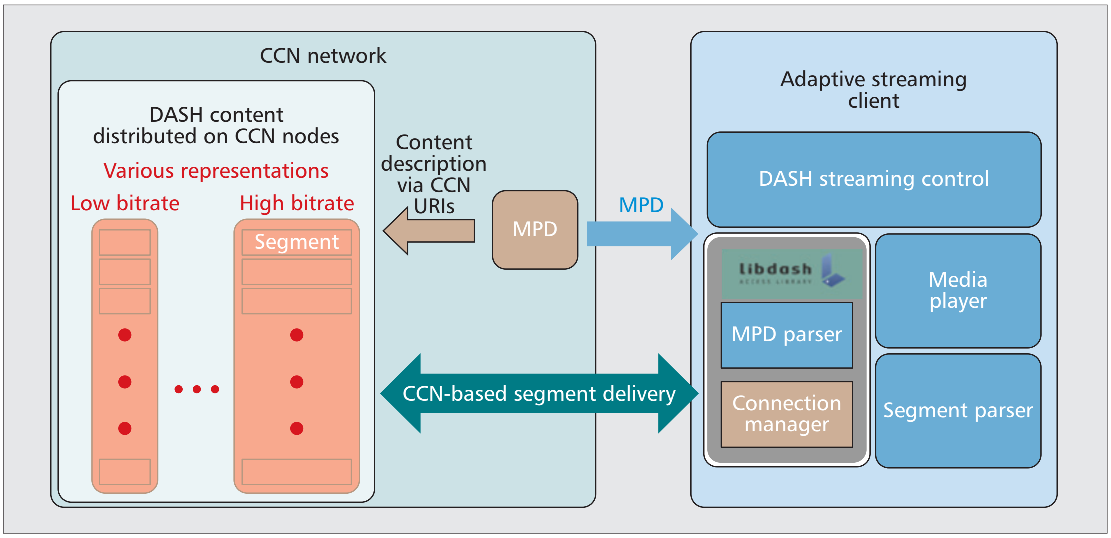

Figure 2. DASH over CCN architecture.

DASH over CCN Architecture

Figure 2 presents the proposed architecture of DASH over CCN, where DASH-related components are marked in orange, CCN-related components in green, and implementation-depended components in blue. The segments are stored within the network and described by an appropriate MPD. The adaptive streaming client is modified with a CCN access module to enable the request and the delivery of DASH segments over CCN. As the delivery of the MPD is not specified in the DASH standard, it is also excluded here. For example, it can be retrieved in the same manner as a segment.

The MPD has to be updated with respect to the segment URIs using a CCN naming scheme instead of HTTP URLs. The excerpt of a simplified MPD is shown in Listing 1 using a custom profile for CCN as indicated in line 1. The *BaseURL* is replaced by a CCNx-compliant URI (line 2) and used as the basis for all URIs of DASH segments. Note that the term BaseURL might be inappropriate as CCN is using URIs instead of URLs, but we did not change the name of the XML element in order to keep it compliant with MPEG-DASH. DASH mandates the concatenation of the BaseURLs on the different hierarchical levels (MPD, Period, Representation) and allows describing the actual segments using different features such as lists, templates, and so on. For example, lines 7 to 10 show a so-called *SegmentList* providing a separate XML element for each segment. In this example, segments of different bitrates are organized in separate directories and after resolving the segment URI with the base URL, the first segment of the 100 kb/s representation would be */itec1/dash/bunny/bunny_2s_100kbit/bunny_2s1.m4s* (i.e. line 2 + line 8).

```
1 <MPD profiles="urn:mpeg:dash:profile:ccn:2013"…>
2 <BaseURL>/itec1/dash/bunny/</BaseURL>
3 <Period start="PT0S">
4   <AdaptationSet bitstreamSwitching="true"
        mimeType="video/mp4">
5     <Representation id="0" bandwidth="101492">
6       <SegmentList duration="2">
7         <Initialization sourceURL="bunny_2s_100kbit/init.mp4"/>
8         <SegmentURL media="bunny_2s_100kbit/bunny_2s1.m4s"/>
9         <SegmentURL media="bunny_2s_100kbit/bunny_2s2.m4s"/>
10        <SegmentURL media="bunny_2s_100kbit/bunny_2s3.m4s"/>
11        <!-- further segments -->
12        <SegmentList/>
13      <Representation/>
14      <!-- further representations -->
15    <AdaptationSet/>
16    <!-- further adaptation sets -->
17  <Period/>
18  <!-- further periods -->
19 <MPD/>
```

Listing 1. Example MPD containing CCN URIs.

The MPD parser processes the MPD and resolves the BaseURLs, which allows the adaptive streaming client to request segments according to the *DASH streaming control*. The DASH streaming control implements the adaptive behavior, which targets a smooth streaming session with short startup delays and high Quality of Experience (QoE) without stalling/re-buffering events. However, this behavior is not defined within the MPEG-DASH standard and is left open to research and industry competition. It may be tailored to specific deployment characteristics, such as in this case for ICN and in particular its CCN and CCNx implementation. The ICN-specific streaming control instructs the *CCN access* module to issue CCN interest packets, which are then forwarded to the appropriate interface(s) based on the information within the FIB. These interest packets are handled within the CCN like any other interest packet leveraging the efficient interest aggregation for popular content as well as the implicit support for multicast. Data packets satisfying the interest packets containing the actual DASH segments are delivered to the client. These data packets are stored on the origin server or any CCN node and with an increasing popularity of the content, these data packets will be distributed across the network resulting in lower transmission delays and reduced bandwidth requirements for the origin servers. Consequently, this reduces the traffic burden for content providers as well as ISPs. Finally, live video streaming is intrinsically enabled thanks to the implicit multicast support.

At the client, the segments within the data packets are processed by the *segment parser* and decoded/rendered by the *media player*, respectively.

Name	Average bitrate	Average switches	Average unsmoothness
Unit	[kbps]	[# of switches]	[Seconds]
Microsoft	1522	51	0
Adobe	1239	97	64
Apple	1162	7	0
DASH VLC	1464	166	0
DASH over CCN	**1326**	**160**	**0**

Table 1. Comparison of DASH over CCN with existing HTTP streaming systems.

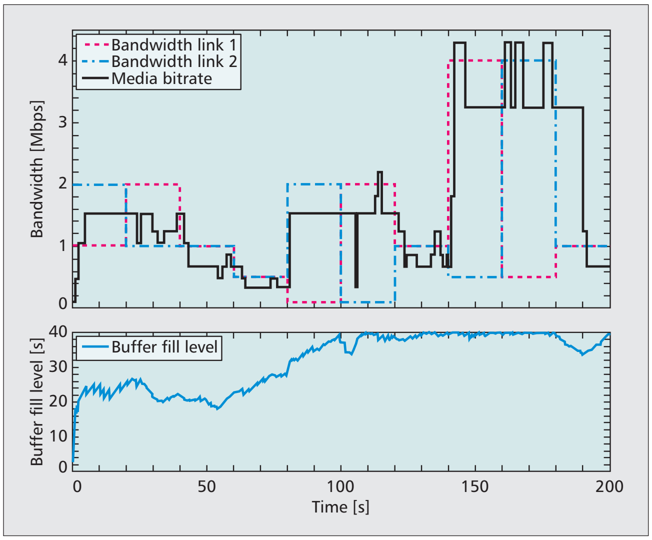

Figure 3. Evaluation results of DASH over CCN with multiple interfaces/links [12].

Open Source Tools, Testbed, and Dataset

For the evaluations of DASH over CCN, we provide various open source tools, a testbed, and a dataset. In particular, we provide two client player implementations: a libdash extension for DASH over CCN, and a VLC plugin implementing DASH over CCN. For both implementations we have used CCNx as a basis.

The general architecture of libdash is depicted in Fig. 2, showing that the library implements an *MPD parser* and an extensible *connection manager*. The library provides object-oriented interfaces for these modules to access the MPD and the available segments. These components are extended to support DASH over CCN and are available in a separate development branch of the GitHub project available at http://www.github.com/bitmovin/libdash. libdash comes together with a fully featured DASH player with a QT-based front end, demonstrating the usage of libdash and providing a scientific evaluation platform. As an alternative, we provide patches for the DASH plugin of the VLC player. These patches can be applied to the latest source code checkout of VLC, resulting in a DASH over CCN-enabled VLC player.

Finally, we provide a DASH over CCN dataset in the form of a CCNx repository. It includes 15 different quality representations of the Big Buck Bunny sequence, ranging from 100 kb/s to 4500 kb/s. The content is split into segments of two seconds, and described by an associated MPD using the presented naming scheme above. This repository can be downloaded from our Web site and is also provided through a publicly accessible CCNx node. Associated routing commands for the CCNx namespaces of the content are provided via scripts coming together with the dataset, which can be used as a public testbed. The client implementations as well as the dataset are publicly available on our DASH research Web site accessible at http://dash.itec.aau.at.

Evaluations

Different evaluations of multimedia transport in ICN have been conducted in the past, including real-time audio conferencing in [8] and live streaming of multimedia content in [9], showing promising results and providing a proof of concept that ICN is suitable for this type of data transfer. Effects of adaptive streaming in ICN, such as the influence of content caching to the streaming performance, have been investigated by [10], showing that caching may cause oscillation of the quality due to wrong bandwidth predictions. Finally, the proposed DASH over CCN architecture has been evaluated in different experiments such as in [11, 12], ranging from protocol overhead and streaming performance evaluations (i.e. comparing to DASH using HTTP) to caching and mobile streaming experiments.

The focus of our evaluations was on the client performance regarding the clients' media throughput compared to existing HTTP streaming systems, and when using multiple links. Therefore, we have created an experimental setup comprising servers and clients including means for network emulation and bandwidth shaping. The details of this experimental setup can be found in [11, 12].

Table 1 shows the results of the evaluation of our DASH over CCN implementation compared to existing HTTP streaming systems. In terms of *average bitrate* received at the client, DASH over CCN is able to compete with existing HTTP streaming systems although it does not outperform all of them. The *number of average switches* during the streaming session is relatively high for DASH over CCN, indicating that the adaptation logic needs more adjustments with respect to characteristics of an information centric network. However, the main goal of adaptive multimedia streaming was achieved as the *number of unsmooth seconds* is zero and, thus, no stalling occurred.

ICN in general, as well as CCN, enable the retrieval of data over multiple links in parallel, and its benefits in combination with MPEG-DASH have been shown in [12]. In particular, it offers the possibility to dynamically switch between the available links depending on their bandwidth capabilities, transparent to the actual DASH client. The result of this evaluation is shown in Fig. 3, with two available *bandwidth links 1* and *2* configured with a pre-defined bandwidth trace that is changing every 20 seconds. The *media bitrate* indicates the throughput achieved at the client and, as shown by the results, the CCN strategy layer always selects the fastest link to retrieve the requested data. For example, in second 20, the bandwidth of link 1 increases from 1 Mb/s to 2 Mb/s and link 2 decreases from 2 Mb/s to 1 Mb/s without affecting the media bitrate. However, the current implementation of the CCN strategy layer causes a notable delay when switching between links, which impacts the media throughput at the client as seen, for example, in second 25 and 100, respectively. That also impacts the adaptation logic of DASH, resulting in switching to a lower rep-

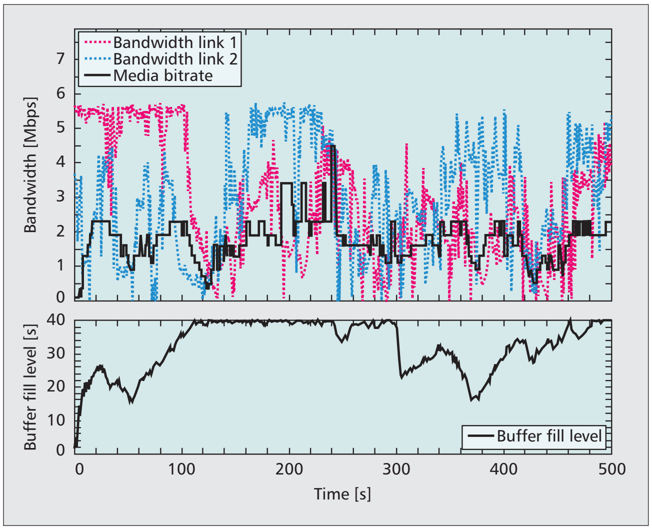

Figure 4. DASH over CCN using multiple links within a mobile/vehicular environment.

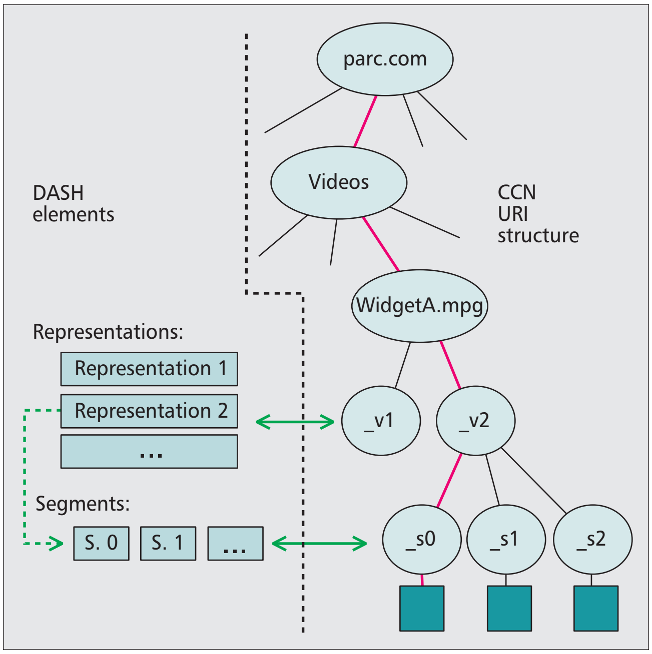

Figure 5. Mapping of DASH representation to the versions of the CCN naming structure [11].

resentation for subsequent segments in order to maintain a smooth playback. Consequently, this shows the benefit of the combination of MPEG-DASH and the multi-link features of CCN, which is also confirmed by the *buffer fill level* indicating no stalls.

In order to confirm our findings we have performed the same experiment using real-world bandwidth traces within a mobile — vehicular — environment. The results are shown in Fig. 4 with an average media bitrate of ~1710 kb/s, which is ~29 percent and ~15 percent higher than if only using bandwidth link 1 and 2 separately. Also, the buffer fill level indicates no stalls, and thus guarantees a smooth streaming throughout the entire session.

The conducted multi-link experiments show both shortcomings and future improvements of the current CCNx implementation. The bandwidth available to the DASH client is only the bandwidth of the fastest link rather than the sum of the bandwidth of all links. That is, the current implementation of CCNx it is not able to split up the data transmission to multiple links, such that the client can leverage the total available bandwidth of all underlying links, which leads us to future research challenges in this emerging area.

Research Challenges

The term *Future Internet* covers a broad field of research activities [3] including efforts within standards developing organizations such as the IETF and its ICN research group (ICNRG), which defines research challenges in this area [13]. In contrast to ICN, adaptive streaming technologies and standards like MPEG-DASH are already well advanced, and large-scale deployments of proprietary media streaming solutions are available by major industry players. Combining adaptive multimedia streaming and ICN raises new challenges, which must be considered in the future, and are highlighted in the following.

Both concepts offer the possibility to request the content from multiple sources. In adaptive multimedia streaming, the available sources are either assigned via an intelligent DNS resolution, or signaled via the manifest, and thus known to the client. For the latter, the adaptation logic at the client may choose the source from which to download segments. In both cases the client is aware of the actual source (or origin server) using its IP address. In ICN, the client does not have the knowledge from which source the requested content is actually served or how many origin servers of the content are available, as this is transparent and depends on the name-based routing. This introduces the challenge that the adaptation logic of the adaptive streaming client is not aware of the event when the ICN routing decides to switch to a different origin server. As most algorithms implementing the adaptation logic are using bandwidth measurements and related heuristics, the adaptation decisions are no longer valid when changing origin servers and potentially cause playback interruptions and, consequently, stalling.

ICN supports the usage of multiple interfaces and a seamless handover between them, which again comes together with bandwidth fluctuations, for example, switching between fixed and wireless, 3G/4G and WiFi networks, and so on. Considering these characteristics of ICN, adaptation algorithms merely based on bandwidth measurements are not appropriate anymore, as potentially each segment can be transferred from another ICN node or interface, all with different bandwidth conditions. Thus, adaptation algorithms taking into account these intrinsic characteristics of ICN are preferred over algorithms based on mere bandwidth measurements.

Another open issue arises as DASH is intended for the adaptive delivery of multimedia content via HTTP. That is, DASH mandates the usage of HTTP URLs for segments within the MPD. However, DASH also enables the usage of profiles, and hence this issue can be solved by introducing a DASH over ICN profile adopting a common ICN naming scheme as shown in the previous section. Therefore, it is required to define such a naming scheme, for example, within the IETF, and its usage in combination with DASH as part of MPEG.

The naming scheme of the segments can be further extended to reflect intrinsic features of ICN concepts like CCN, e.g. versioning and segmentation support as proposed in [11]. As segmentation is already compulsory for multimedia streaming, it can be leveraged for DASH-based streaming over CCN. The CCN versioning can be adapted to signal different DASH representations as depicted in Fig. 5. That is, *representations 1* and *2* are mapped to versions of the CCN URI [4] indicated with the suffix *_v1* and *_v2*. This also applies for the segmentation support of ICN, which can be adopted to reflect the segment structure of the DASH content. This may be a good approach for a common DASH over ICN naming scheme, but also enables further possibilities such as implicit adaptation of the requested content to the clients' bandwidth conditions on a network layer, rather than on the application layer, enabling an adaptation on a hop-by-hop basis instead of on a client-server basis. As an extension to this, DASH segments could be aggregated automatically by the CCN nodes if bandwidth conditions of the corresponding interfaces or routing paths allow it, which would reduce the amount of interest packets needed to request the content. However, such approaches need further research, specifically in terms of additional adaptation mechanisms and processing power needed at the CCN nodes.

Conclusions

Adaptive multimedia streaming over HTTP is currently widely accepted as the main delivery mechanism of multimedia content in the Internet. MPEG-DASH is the result of various proprietary formats developed by the industry, and research focuses on the optimization thereof in terms of efficient delivery and Quality of Experience. ICN is in its infancy and mostly driven by research institutions but, if widely adopted, is a promising candidate for the efficient delivery of multimedia content in the Future Internet.

DASH and the underlying technology is actually agnostic to the delivery infrastructure, which currently is limited to HTTP for which it scales well but it can be easily extended to other means of transport such as ICN. In this article we have shown that DASH and ICN share some basic concepts and an effective integration of both technologies can be achieved, but also introduces new challenges that need further attention in the future.

Acknowledgment

This work was supported in part by the EC in the context of the ALICANTE (FP7-ICT-248652), SocialSensor (FP7-ICT-287975), and QUALINET (COST IC 1003) projects, and was partly performed in the Lakeside Labs research cluster at AAU.

References

[1] Sandvine, "Global Internet Phenomena Report 1H 2013," *Sandvine Intelligent Broadband Networks*, 2013.

[2] I. Sodagar, "The MPEG-DASH Standard for Multimedia Streaming over the Internet," *IEEE MultiMedia*, vol. 18, no. 4, 2011, pp. 62–67.

[3] J. Pan, S. Paul, and R. Jain, "A Survey of the Research on Future Internet Architectures," *IEEE Commun. Mag.*, vol. 49, no. 7, 2011, pp. 26–36.

[4] V. Jacobson *et al.*, "Networking Named Content," *Proc. 5th Int'l. Conf. Emerging Netw. Experiments and Technologies (CoNEXT '09)*, 2009.

[5] P. Frojdh *et al.*, "Adaptive Streaming Within the 3GPP Packet-Switched Streaming Service," *IEEE Network*, vol. 20, no. 2, Mar.-Apr.2006, pp. 34–40.

[6] L. De Cicco and S. Mascolo, "An Experimental Investigation of the Akamai Adaptive Video Streaming," *Proc. 6th Int'l. Conf. HCI in Work and Learning, Life and Leisure: Workgroup Human-Computer Interaction and Usability Engineering (USAB'10)*. Berlin, Heidelberg: Springer-Verlag, 2010, pp. 447–64.

[7] A. Detti *et al.*, "Offloading Cellular Networks with Information-Centric Networking: The Case of Video Streaming," *Proc. Int'l. Symp. on a World of Wireless, Mobile and Multimedia Networks (WoWMoM 12)*, IEEE, San Francisco, CA, 2012, pp. 1-3.

[8] V. Jacobson *et al.*, "VoCCN: Voice-over Content-Centric Networks," *Proc. 2009 Workshop on Re-Architecting the Internet (ReArch '09)*, 2009, pp. 1-6.

[9] Z. Chen, R. Chen and J. Cao, "Live Streaming with Content Centric Networking," *Proc. 3rd Int'l. Conf. Netw. and Distributed Computing*, Hangzhou, China, 2012.

[10] R. Grandl, K. Su, and C. Westphal, "On the Interaction of Adaptive Video Streaming with Content-Centric Networking," *Proc. 20th Packet Video Workshop 2013*, San Jose, CA, Dec. 2013.

[11] S. Lederer *et al.*, "An Experimental Analysis of Dynamic Adaptive Streaming over HTTP in Content Centric Networks," *Proc. IEEE Int'l. Conf. Multimedia and Expo 2013*, San Jose, CA, July 2013.

[12] S. Lederer *et al.*, "Adaptive Streaming over Content Centric Networks in *ICC*, Budapest, Hungary, June 2013.

[13] D. Kutscher *et al.*, "ICN Research Challenges," *IETF Information-Centric Networking Research Group (ICNRG)*, URL: http://datatracker.ietf.org/doc/draft-kutscher-icnrg-challenges (visited: 01.10.2013).

Biographies

STEFAN LEDERER is an assistant professor at the Institute of Information Technology (ITEC), Multimedia Communication Group, and CEO at bitmovin GmbH. He received his M.Sc. (Dipl.-Ing.) in computer science and M.Sc. (Mag.) in business administration from Alpen-Adria-Universität (AAU) Klagenfurt. His research topics include transport of modern/rich media, multimedia adaptation, QoS/QoE, and future Internet architectures, about which he published more than 15 papers. He participated in several EC-funded Projects (ALICANTE, SocialSensor) and in the MPEG-DASH standardization, where he contributed several open source tools (DASHEncoder, DASHoverCCN VLC plugin) and datasets.

CHRISTOPHER MÜLLER is an assistant professor at the Institute of Information Technology (ITEC), Multimedia Communication Group, and CTO at bitmovin GmbH. He received his M.Sc. (Dipl.-Ing.) from the AAU Klagenfurt. His research interests are multimedia streaming, networking, and multimedia adaptation, about which he published more than 20 papers. He gained practical expertise in various companies (Infineon, Dolby Laboratories Inc. LA, etc.) and participated in the MPEG-DASH standardization, contributed several open source tools (VLC plugin, libdash), and participated in several EC-funded projects (ALICANTE, SocialSensor).

CHRISTIAN TIMMERER is an associate professor in the Institute of Information Technology (ITEC), Alpen-Adria-Universität Klagenfurt, and CIO at bitmovin GmbH. His research interests include immersive multimedia communication, streaming, adaptation, and Quality of Experience. He was the general chair of WIAMIS'08 and QoMEX'13 and has participated in several EC-funded projects, notably DANAE, ENTHRONE, P2P-Next, ALICANTE, and SocialSensor. He also participated in ISO/MPEG work for several years, notably in the area of MPEG-21, MPEG-M, MPEG-V, and MPEG-DASH. He received his Ph.D. in 2006 from the Alpen-Adria-Universität Klagenfurt. Follow him on twitter.com/timse7 and subscribe to his blog (blog.timmerer.com).

HERMANN HELLWAGNER is a full professor of Informatics in the Institute of Information Technology, Alpen-Adria-Universität Klagenfurt, Austria, leading the Multimedia Communications group. His research areas are distributed multimedia systems, multimedia communications, and quality of service. He is a senior member of the IEEE, a member of the ACM, GI, OCG, and vice president of the Austrian Science Fund.