International Journal of Semantic Computing

A hybrid NDN-IP Architecture for Live Video Streaming: From Host-based to Content-based delivery to improve QoE

Ishita Dasgupta

College of Information and Computer Sciences,
University of Massachusetts Amherst, 140 Governors Dr,
Amherst, Massachusetts, 01002 USA ishitadg@cs.umass.edu

Susmit Shannigrahi

Computer Science Department
Tennessee Tech University, 331 Bruner Hall,
Cookeville, Tennessee, 38501, USA sshannigrahi@tntech.edu

Michael Zink

Electrical and Computer Engineering,
University of Massachusetts Amherst,213B Knowles Engineering Building
Amherst, Massachusetts, 01002 USA zink@ecs.umass.edu

With live video streaming becoming accessible in various applications on all client platforms, it is imperative to create a seamless and efficient distribution system that is flexible enough to choose from multiple Internet architectures best suited for video streaming (live, on-demand, AR). In this article, we highlight the benefits of such a hybrid system for live video streaming as well as present a detailed analysis with the goal to provide a high quality of experience (QoE) for the viewer. For our hybrid architecture, video streaming is supported simultaneously over TCP/IP and Named Data Networking (NDN)-based architecture via operating system and networking virtualization techniques to design a flexible system that utilizes the benefits of these varying Internet architectures. Also, to relieve users from the burden of installing a new protocol stack (in the case of NDN) on their devices, we developed a lightweight solution in the form of a container that includes the network stack as well as the streaming application. At the client, the required Internet architecture (TCP/IP versus NDN) can be selected in a transparent and adaptive manner.

Based on a prototype we have designed and implemented maintaining efficient use of network resources, we demonstrate that in the case of live streaming, NDN achieves better QoE per client than IP and can also utilize higher than allocated bandwidth through in-network caching. Even without caching, as opposed to IP-only, our hybrid setup achieves better average bitrate and better perceived visual quality (computed via VMAF metric) over live video streaming services. Furthermore, we present detailed analysis on ways adaptive video streaming with NDN can be further improved with respect to QoE.

Keywords: Named Data Networking, Live Streaming, Software Defined Networking, quality of experience, VMAF, HEVC

1. Introduction

With the advent of applications like Twitch, Facebook Live, live streaming has increased tremedously in popularity occupying upto 17% of the video traffic by 2022 [45]. Therefore mechanisms are required that deliver efficient Quality of Experience (QoE) for live video streaming over existing internet infrastructures.

Apart from the traditional TCP/IP, multiple internet architectures have been proposed for the future internet [1][2][39]. Information Centric Networking (ICN) [1] is one such future internet architecture that has the goal of replacing the host-centric approach of the current (TCP/IP-based) Internet with a content-centric approach. Named Data Networking (NDN) [49] is an instantiation of ICN that identifies and serves data by name instead of their location. In NDN, the communication is client-driven in the sense that it sends out "Interest packets" requesting content and receives a response as "Data packet" from the producer. Both these packets are stored in the NDN router for some time in Pending Interest Table (PIT) and Content Store cache, respectively. NDN routers forward packets using information from their Forwarding Information Base (FIB) and if an Interest arrives for a Data packet that is stored in the router's Content Store, the request is directly served from the cache. NDN can improve video streaming application performance due to its in-network caching and multi-path forwarding capabilities [36]. Obviously, the in-network caching characteristics of NDN are well-suited for live streaming events like the Superbowl or the FIFA Worldcup final, where the same content is requested by many viewers simultaneously (potentially in geographic proximity). Since NDN requires the replacement of the Layer 3 protocol, the immediate widespread adoption of this new approach is difficult. Experiences with the long and painful rollout of IPv6 and some new TCP flavors, which all require changes in the operating systems of virtually all nodes (routers and hosts) in the Internet, have clearly shown how cumbersome this process can be as well.

Motivated by these past experiences, we design, implement, and evaluate an approach that can transparently choose and adapt to multiple internet architectures (TCP/IP and/or NDN). This requires a portable, flexible and an easy to configure system that can be easily deployed on a variety of hardware platforms and dynamically adapt to application and network demands. To achieve this goal, our approach employs Software-Defined Networking (SDN) and Network Function Virtualization (NFV) [6] for network virtualization. Further, our approach is based on operating system virtualization as offered by Platform-as-a-Service container systems like Docker [3] or Singularity [19]. Such an approach bears the benefit that a particular application can run in a container with a completely pre-configured environment that does not require any user administration. Also, this bypasses the requirement for widespread replacement of the Layer 3 protocol in the case of NDN. Previous work [44] focused on the implementation of network elements that could support alternative internet protocol stacks parallelly, allowing the clients to stream videos over both TCP/IP and NDN. However, our work presented in this article focuses

on the following new contributions:
1. Evaluating our Hybrid architecture with HEVC dataset over a multi-tier network infrastructure: We analyze our hybrid infrastructure in comparison to traditional IP with respect to live video streaming experience. This end-to-end approach supports transparent architecture selection all the way to the application running on end systems.
2. Newer Detailed QoE analysis on live streaming with NDN & IP: In live streaming, quality of the video playback is as important as the quality of the video. Hence with the focus to optimize user's overall QoE, we inspect the following:

- Predict the viewability of the video content streamed at the client using Video Multimethod Assessment Fusion (VMAF) metric. This metric helps us evaluate the performance of our setup and live streaming sessions with respect to a user's perception.
- Average bitrate & bandwidth utilization by NDN and IP clients.
- The effect of caching on live streaming QoE.
- Existing drawbacks due to the way ABR algorithms work with NDN on user's QoE.
- Performance evaluation of NDN vs. IP clients across server and client-side bottlenecks.

To summarize, this article proposes a hybrid setup that is flexible to the network as well as the underlying physical layer such that it simultaneously supports multiple Internet architectures for better QoE in an ABR application. Furthermore, with our detailed QoE analysis, we justify using NDN along with our hybrid setup to improve live video streaming experience with minimal end-user involvement.

2. SDN and NFV Support for Parallel Live Video Streaming

The inherent caching properties of NDN have proven to support live streaming applications well [44]. However, to benefit from this improvement, significant changes throughout the network have to be performed to support NDN, including installation at end systems. Hence, we propose a hybrid approach where clients can stream videos over both IP and NDN without needing NDN kernel support on hosts or fully replacing the TCP/IP protocol stack with NDN. We achieve this by utilizing virtualization techniques for standalone NDN-based container applications that can run on any client host supporting containerization. Further to enable the parallel support of IP and NDN protocol stacks, we combine SDN and NFV. Our primary goal is to build and evaluate an environment that can utilize the benefits of both these network architectures simultaneously in the domain of video-streaming while optimizing resource use and user experience. In that context, we present a brief overview on Live streaming with NDN and network components that enable flexible hybridity of our architecture.

2.1. *ABR and DASH*

Virtually all commercial streaming services that offer either VoD, live streaming, or both, use ABR streaming. With ABR, the streaming rate, and thus the resulting video quality, can be adapted to the available bandwidth between the video server and the client. One of the most popular realizations of ABR streaming is the MPEG's Dynamic Streaming over HTTP (DASH) standard [41]. Its popularity can mostly be attributed to the facts that *i)* DASH-format videos can be streamed from any kind of HTTP server; *ii)* adaptation logic resides in the client, which makes DASH highly scalable; and *iii)* it is an open standard. In addition to DASH, there exist proprietary ABR implementations like Microsoft's SmoothStreaming [26], Apple's HTTP Live Streaming (HLS) [14], and Adobe's HDS [15] exist. The work presented in this article is based on the open DASH standard.

All ABR approaches mentioned above rely on the same principle, which divides a video into segments (usually between 1 - 10 seconds in duration). Each segment is then encoded in different quality versions, which require different bandwidths to be streamed in real-time from the server to the client. For each video, a media presentation description (MPD) file is created that contains meta-information about the segments and the different qualities they are encoded in. A streaming session starts with the client retrieving the MPD file. In the case of live streaming, a dynamic version of this MPD is required such that the client receives a new version of the MPD file once additional segments of the live video become available at the source. The playback bitrate of the next segment to be retrieved is calculated using the information given by the MPD file, the segment download rate determined at the client, as well as the actual fill of the video player buffer. As opposed to preprocessing the entire video beforehand in video-on-demand services, only small chunks of the newly arriving feed can be segmented and encoded in different qualities in the case of live streaming.

2.2. *NDN and Live Streaming*

NDN provides a set of features that are well-tailored for live streaming. Especially, the feature to cache content at each router (even if only small amounts can be cached compared to a CDN or web cache) supports live streaming where the same content is requested simultaneously by many viewers that might have a high degree of geographical locality. This creates an efficient multicast mechanism. Thus, many user requests can be served from a router in the network instead of the origin server or a CDN edge server. Opposed to the traditional IP case where content sources are less diversified, in the case of NDN, DASH segments may come from various sources (a data publisher that is closest to the consumer, an in-network cache, and so on), when multiple sources are available. Additionally, each DASH segment will be chunked into multiple NDN Data packets that are 8800 Bytes in size (by default, the size of the packets is configurable). These Data packets may come from various sources as well. Further research is needed to create congestion control algorithms

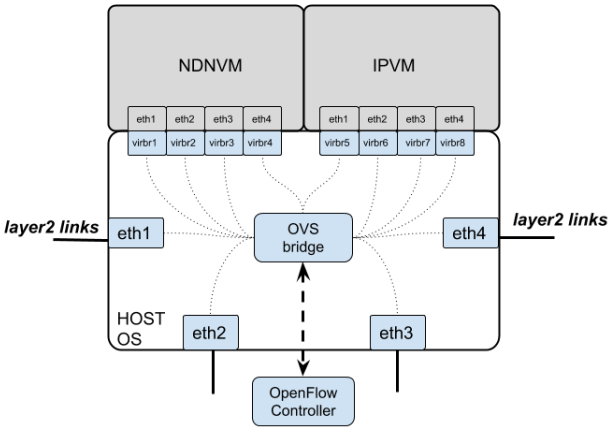

Fig. 1: SDN-NFV setup in an intermediate router

that can take into account this heterogeneity of Data sources in NDN [35].

The advantages of NDN for efficient delivery of video streams have been investigated in works such as [11] and [37], which make ICN-based protocols a likely candidate for the transport of live video. While the general feasibility of supporting live streaming with NDN has been demonstrated earlier [44], the approach in this article focuses on providing transparent solutions that do not require users' involvement to configure the underlying network technology. Additionally, this article provides a much more detailed evaluation of live video streaming based on our hybrid streaming architecture.

2.3. *Network Devices*

We have chosen an SDN supported NFV approach for the network device design that handles IP and NDN traffic in an isolated and adaptive fashion. This architecture enables flexible topology setup and efficient resource use. In addition, it can be easily extended to support other network layer protocols (e.g., IPv6 or IPSec). In our particular use case, this architecture is used to dynamically configure customized hybrid routers that simultaneously support IP(v4) and NDN. In this architecture, SDN is used to internally (within the device) direct traffic from the physical network interface to the respective virtual interface of the respective router. This is realized by running openvswitch [32] (OVS) on the Host OS of the node. An SDN controller implements rules for isolating NDN from IP traffic in the OVS-based SDN switch as shown in Figure 1. While we use a centralized OpenFlow controller in our prototype, the architecture is designed such that any type of controller architecture (centralized or distributed) can be supported. Furthermore, our architecture reduces resource utilization in the case of live-streaming as we do not need client machines with

NDN-supported kernels, only hosts that can run NDN-based container applications whenever live-streaming is requested.

3. Evaluation Setup

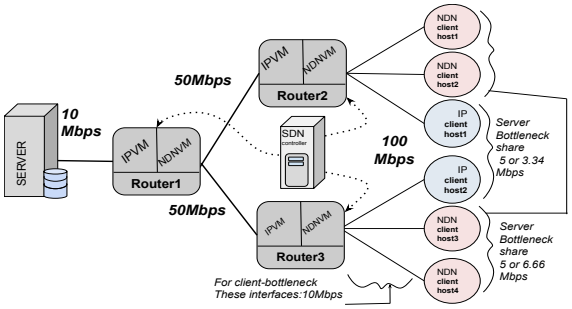

Fig. 2: Topology for our Evaluation setup

3.1. *Experimental Setup*

In this section, we describe the components of the architecture we employ for the evaluation of our approach.

Topology: Figure 2 shows the topology used for the evaluation experiments, which consists of a server, three intermediate routers (OpenFlow enabled), six Ubuntu clients, and a centralized SDN controller. The server node contains the videos being served for live streaming. The clients can be classified as an NDN client or IP client depending on which Internet architecture is being used by the client node for streaming. We chose this topology as we wanted to evaluate tiered cache effects amongst NDN clients in comparison with traditional IP clients.

We implemented this topology on the Cloudlab testbed [7] which offers bare metal servers providing flexibility to configure individual nodes exactly to our needs. CloudLab supports reproducibility, allowing us to share the complete setup and make the execution and outcome of our evaluation repeatable by other researchers. The decision to use a testbed for our evaluation instead of using a simulation or emulation environment like ndnSIM [23] or Mininet [13], respectively, came from the consideration that the latter would abstract several important factors that have an impact on the overall performance of our approach.

Network Setup: Figure 2 shows the available bandwidth at each link interface. We employ traffic control (tc) [8] to enforce the maximum bandwidth limits and the share of bandwidth that is available for NDN and IP traffic. Preliminary experiments revealed that TCP saturates the available bandwidth quickly due to highly optimized congestion control, slowing down NDN transfers. On the contrary, NDN is implemented in the application layer and currently lacks sophisticated congestion control mechanisms. If not mentioned otherwise, the link between Server and

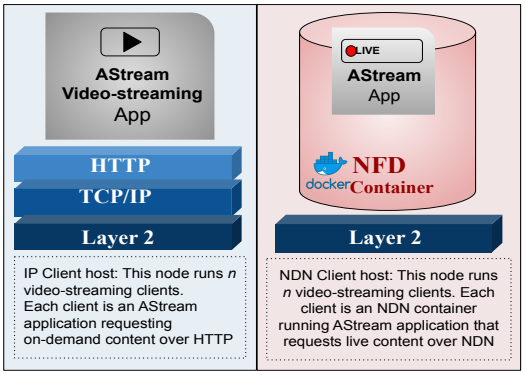

Fig. 3: Execution environment for IP (left) and NDN (right) ABR streaming clients.

Router1 is shared equally between IP and NDN traffic. For the virtualized NDN routers (running on the physical router nodes 1-3 in Figure 2) the cache size was varied between 0MB, 250MB, and 500MB.

As we will see in the case of server-side and client-side bottleneck experiments in Sects. 4.2.1 and 4.2.5, the chosen bottlenecks were 5Mbps to serve upto 20 clients and 10Mbps to serve 5 clients, respectively. In the server-side scenario, this was done to maintain the proportionality of a limited bandwidth to multiple customers. In the case of client-side bottleneck, the highest quality segment in our dataset is 3.8Mbps. To serve 5 clients at highest quality, a bandwidth of 20 Mbps would be required, hence the bottleneck is set to 10 Mbps. In other words, these bottleneck bandwidth numbers were chosen keeping proportionality with real world in mind and not the absolute values themselves.

Virtualization: As introduced in Sect. 2, network virtualization is implemented using a combined NFV-SDN setup, whereas containers are used for application-level virtualization in the case of NDN live streaming clients. The Kernel-based Virtual Machine (KVM) hypervisor is used at the routers to handle IP and NDN traffic separately and in an isolated fashion. Figure 1 shows the virtualized network configuration of Routers 1-3. The physical layer 2 links of the host are mapped to the virtual interfaces of the internal VMs via OpenVswitch [32], which is managed by the SDN OpenFlow controller. Internally in the router, the traffic is routed from the server to client nodes via these virtual VM interfaces at the intermediate routers.

Using this approach, the controller uses the EtherType field (0x8624 for NDN, 0x0800 for IPv4) of an incoming frame to determine which virtual interface it has to be forwarded to. As shown in Figure 1, an incoming IP packet on interface `eth1` would be forwarded to virbr5, while an outgoing NDN packet from virbr4 would be forwarded to `eth4`. More details on this configuration can be found in [44]. At the client side, Docker containers [25] are used for the NDN live streaming application. For the networking aspect of our containers, we used Docker macvlan networks [5] that connect each container's virtual interface to a host's physical interface. This

also lets us monitor and regulate the traffic at Layer 2, which is needed for the NDN clients (see Figure 3). The choice for Macvlan networking bypasses the necessity to make the host entirely visible (as in host networking [4]) or the need to assign individual IP addresses (as in ipvlan networking [5]). User-defined or default bridge networking on the docker containers was not an option because communication between the containers running on the host was not required.

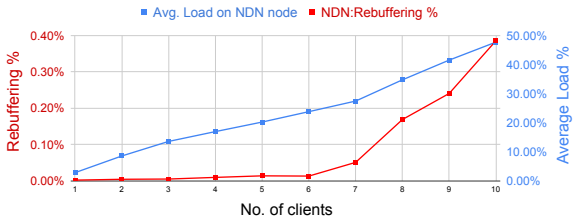

Fig. 4: Effect of increasing clients on CPU Load and Rebuffering percentage.

Deciding the number of NDN containers per host: For the large scale experiment scenarios we describe in Sect.4, using one physical node per client would have required a total of 65 physical nodes. Due to resource limitations of the CloudLab testbed, using such a large number of physical nodes is not feasible. Therefore, we had to resort to run several NDN containers on the same physical nodes. We ran experiments to decide how many clients we could safely run on each physical node. We ran a series of experiments where we constantly increase the number of containers (NDN) running on a physical node. We start with a single client streaming over NDN and monitor the CPU load during the process as well as the resulting Rebuffering percentage (see Sect. 3.3 for the definition of the metric). This experiment is then scaled up to ten clients and we observe an increase in the Rebuffering percentage along with CPU load, as shown in Fig. 4. While both metrics increase proportionally to the number of clients, the Rebuffering percentage for the case of five NDN container-based clients per physical node is only 0.01%. The other QoE metrics (e.g., Spectrum and QS) show the same behavior though we do not include them for brevity. Since five clients do not significantly affect the QoE, we decided to run five NDN containers per physical host for most of our large-scale experiments.

Video: For our evaluation, we made use of the Big Buck Bunny video [33], which is encoded in the DASH/HEVC format and supports up to 10 different quality bitrates. `(0.15 Mbps, 0.28 Mbps, 0.45Mbps, 0.61 Mbps, 0.92 Mbps, 1.54 Mbps, 2.26 Mbps, 2.89 Mbps, 3.4 Mbps, 3.8 Mbps)`. Each quality representation is a 10-minute long video divided into 2-second segments. We use this well-known encoding format for our dataset since it represents the popular, state-of-the-art HEVC coding standard. It should also be noted that the same dataset is used for IP and NDN evaluations, hence the data itself or its encoding type, is of

less importance with the respect to the comparative performance evaluation.

Video Server: Since DASH is employed as the streaming technique in this work, we use a regular web server (vanilla Apache2) for IP-based streaming. For NDN, the live-streaming content needs to be named at the granularity of a DASH segment's quality, which can be achieved with minimal effort. For a given video "v" whose DASH segment "n" is of quality "q", the segment is named as "$<ndn\_prefix>/<v>\_<q>\_<n>.m4s$." This content is served using *ndn-python-repo* [18] which creates a repository at the server containing the DASH video segments in their respective quality levels. Combined with *ndncatchunks* [31] on the client-side, it delivers data based on content names. *ndncatchunks* implements a TCP CUBIC-like [12] congestion control algorithm that can adjust the data transfer rate based on the observed network conditions (e.g., congestion, packet loss).

Video Client: For the streaming client, we use a python-based video streamer, *AStream*[16] that supports multiple DASH adaptation algorithms. Here, *AStream* uses HTTP libraries and *ndn-python-repo* combined with *ndncatchunks* to download video segments over IP and NDN, respectively. We selected BOLA [43] as the bitrate adaptation algorithm. We choose BOLA because it is a near-optimal state-of-the-art adaptive bitrate streaming algorithm for our player at the client and is also a part of the DASH reference player [42]. As long as the same ABR algorithm is used by NDN and IP for comparative analysis of the QoE, we can always use alternate ABR algorithms with this setup in the future.

3.2. *Live Streaming*

Since our focus is only on the streaming part, we ignore the production process of live content and only focus on the content delivery. For the evaluation of our approach, we use the Big Buck Bunny video as described in Sect. 3.1, where only the very first client starts requesting segments from the very beginning of the video. We specify that this first request occurs at t_0. We log this time at the video server and once a request from a new client arrives (e.g. at $t = t_1$), we determine the starting segment for that client as $(t_1 - t_0)/segmentlength$. For example, if the first client starts requesting $t_0 = 0$ seconds and the next client request is received at $t_1 = 10$ seconds and we assume a video segment length of 2 seconds, then the second client is served the video from segment 5 onwards. To allow the client to determine the correct starting segment, this information is transmitted from the server to the client in the dynamic MPD file.

3.3. *Metrics*

Since one of our goals is to optimize QoE, we use the following metrics that are widely accepted as good representations for viewers' perceived quality.

- Average Quality Bitrate (AQB): One of the objectives of quality adaptation algorithms is to maximize the average quality bitrate of the streamed video.

For a comprehensive QoE representation, we need to combine this metric with the *Number of Quality Switches*.

- Number of Quality Switches ($\#QS$): This metric is used together with AQB to draw quantitative conclusions about the perceived quality (QoE). For example, for two streaming sessions having the same AQB, the session with the lower $\#QS$ will be perceived better by the viewer.
- Spectrum (H) [50]: This metric combines the variation of the video quality bitrate around the average bitrate. A lower H indicates better QoE.
- Rebuffering percentage (RB): The average rebuffering percentage is given by:

$$RB = \mathsf{E}\left[\frac{t_a - t_e}{t_e}\right]\%, \tag{1}$$

where t_a is the actual playback time and t_e is the entire video length in seconds.

- VMAF score ($VMAF$): Video Multimethod Assessment Fusion [20] is a metric designed by Netflix to reflect human perception of video streaming quality. We adopted this metric to gain additional assessment of the perceived video quality at the clients. VMAF a score between two videos, original and distorted, indicates how viewable a distored video is in reference to the original. VMAF uses Support Vector Machine regression trained on a Netflix video data set and existing image quality metrics to provide a score in the range of 0-100, where 100 means distorted video has a quality identical to the original video. For example, if the client streams the video at 3.8 Mbps throughout the streaming session of 10 minutes, the VMAF score would ideally be 100. Hence, VMAF is also highly dependent on the quality of reference.

  For each of our streaming sessions, we compare the VMAF scores of the video streamed at the client (where each segment could be requested at a different quality) to the reference video. For all of our evaluations the reference video was set to the highest bitrate and resolution available in the dataset (3.8 Mbps, 1920x1080 pixels). It should be noted that to compute VMAF between videos of different resolution, the lower resolution was upsampled in resolution to match our reference video.

4. Evaluation Results

The aim of this article is to present the feasibility and benefits of our hybrid network model over a traditional IP-based one with respect to live video streaming. Additionally, we motivate why NDN should be favored over IP in live streaming scenarios. In this light, the results show that our hybrid model achieves higher bitrate as NDN achieves overall higher bandwidth utilization than traditional TCP/IP. In this article, we also compare the perceived video quality of NDN and IP clients

by computing VMAF metrics and found that NDN always outperforms IP across different scenarios. We observe that improvement in certain QoE metrics for NDN live streaming clients increases with increasing cache size but with diminishing improvements. Finally, we identify drawbacks of implementing adaptive streaming with NDN that lead to oscillation effects. Based on our findings we suggest changes in the streaming approach to further improve the QoE even with a client-side bottleneck. To evaluate our approach, we carried out small as well as large-scale experiments on the hybrid model that also tests the scalability and robustness of our approach.

It should be noted that for each experiment, all clients start streaming at the same time and the reported results were accumulated from an average of 10 streaming sessions.

Experiment I: Small-Scale This scenario serves as a baseline in observing the impact of increasing clients (running on a single physical node) on the overall QoE. For this initial small-scale evaluation of our approach, we run one NDN streaming client application in a docker container on each of the four physical client nodes (one container per node as shown in Fig. 2), while two IP-based streaming applications run on each of the two physical client nodes (two per node).

Experiment II: Large-Scale For this set of large-scale experiments, we evaluate the following setups:

a) 20 NDN & 20 IP: Here, we increase the number of clients to 20. Since we cannot increase the number of physical nodes in the topology, we have to run 20 IP clients on two physical nodes (10 on each) and 5 NDN containers on four physical nodes, each (see Fig.2). Apart from the QoE analysis, we chose this setting to also study **i)** the effect of varying cache sizes (0MB, 250MB, 500MB) on the performance of NDN clients and **ii)** compare the performance of NDN and IP clients when the bottleneck is on the client-side instead of server-side.

b) 40 NDN & 20 IP: Further testing the scalability, the number of NDN clients is increased to 40. This increase is motivated by the assumption that popular live streaming events will be watched by many viewers (almost in a flash crowd style) putting additional stress on the system. To adjust for the imbalance between the number of NDN and IP clients, we adjust the bandwidth allocation on the `10Mbps` link between the server and Router 1. As explained in 3.1, this number was chosen to study the effect of limited server-side bandwidth on midgress traffic. For this experiment, we allocate $2/3rd$ of the bandwidth to NDN sessions and $1/3rd$ to IP sessions (proportional to the number of clients of each type).

The average QoE metrics for all experiments are shown in Table 2. In this table, we present the Rebuffering percentage, #QS, average bitrate, spectrum, and VMAF reported across 10 live streaming sessions for experiments I, IIa & IIb. In optimizing the viewer's QoE, a higher avg. bitrate, VMAF and lower Rebuffering percentage, #QS, spectrum is preferred. A more detailed analysis on these results is presented in the following sections. The CDFs for QoE metrics for Experiments I, IIa & IIb are presented in Fig. 5, 6, 7 respectively.

With respect to the end-to-end video-streaming architecture, this work focuses on the subset of this architecture from post video-processing to the distribution to the viewers. Henceforth, we do not regard the latency from video capture (camera or camera set at a live event) to making segments available on the video server in our evaluation. To estimate latency differences between IP and NDN clients between the central server and end-client, we compare the download time differences for identical content with Experiment IIa 's 500MB cache setup. Our results show that IP clients take an average of 0.5 seconds more time than NDN clients to download the same content. In-network caching in case of the NDN clients lead to lower download times, hence, reduced average latency.

4.1. *Hybrid Streaming Architecture Analysis*

Our hybrid model has the flexibility to stream videos over multiple network architectures without affecting each other adversely. On top of providing this benefit, the overall impact on QoE also needs to be investigated. This evaluation allows the comparison between an IP-only scenario and a mixed NDN/IP scenario (as presented in experiments I & II). For the IP-only case, we run 40 IP live streaming clients. Similar to NDN clients in the hybrid setup, half of the 40 IP live streaming clients also run in a container. Table 1 shows the average QoE metrics for the IP-only case and the NDN/IP case. For a fair comparison with cache-less IP, the cache size at the NDN routers was set to `0MB` in this experiment. Table 1 shows that NDN live streaming performs better with respect to average bitrate, even if no caching is performed. In NDN, requests for the same content that are close in time are aggregated in the Pending Interest Table (PIT). When corresponding data arrive at an NDN router it will be multicast to more than one client if more than one outgoing interface is registered in the PIT, thus, creating an inherent multicast mechanism. We attribute the improved performance to this inherent multicast characteristics that complement live streaming scenarios well. In terms of perceptual quality of the video, it is seen that our hybrid setup gives a better VMAF score than its IP-only counterpart. In other words, when compared to the best quality video available in our dataset, clients in our hybrid setup streamed a better quality video as opposed to when clients only use IP. However, this bitrate and VMAF enhancement is accompanied with increased #QS, Rebuffering Ratio and Spectrum. We identify the cause of this increase and present a solution to this drawback in the next section.

4.2. *Live Streaming with NDN vs. IP*

This section first analyzes the benefits of using NDN over IP for live video streaming due to its in-network caching capabilities. We further study the effect of increasing the cache size of NDN routers with respect to QoE. After discussing the benefits, we analyze the drawbacks of using NDN in relation to live video streaming and suggest possible solutions. Finally, we show that with our suggested improvements,

Table 1: Average QoE for IP-only and NDN/IP case.

Scenario	Rebuf %	#QS	Avg. bitrate	Spectrum	VMAF
20 IP, 20 IP, 50/50BW split					
IP	0.72	33.43	0.25	23.69	42.34
20 NDN, 20 IP, 50/50 BW split, 0MB Cache					
NDN-IP Hybrid (0-cache)	3.08	68.46	0.32	95.01	55.81

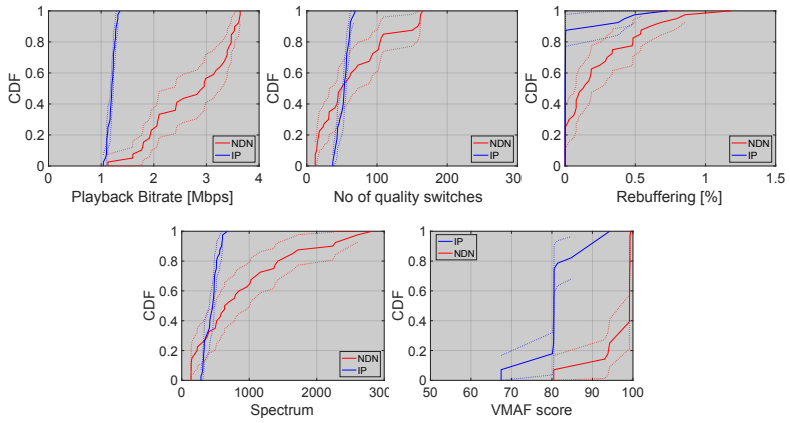

Fig. 5: Cumulative Distribution Functions for QoE metrics for the case of 4 NDN and 4 IP clients.

NDN outperforms IP across all QoE metrics with server-side as well as client-side bottlenecks.

4.2.1. *Higher bitrate & bandwidth utilization by NDN*

NDN clients report higher average playback bitrate across all experiments as can be observed from the corresponding column in Table 2 and their respective CDFs (Figures 5, 6, 7). While throughput for NDN traffic is limited to `5Mbps` (50% of the `10Mbps` link between server and Router 1) in experiments I & IIa, the combined average bitrate is almost double (4*2.74=`11Mbps`) or higher (20 * 0.75=`15Mbps`). Similar conclusions can be drawn from experiment IIb results, as cumulative avg. bitrate increases with increasing clients. This confirms our hypothesis that NDN benefits live streaming scenarios due to its inherent in-network caching and improved handling of potential NDN segment retransmissions. In comparison, the cumulative average bandwidth for IP clients for both experiments stays slightly less or equal to (4*1.2=`4.8Mbps`, 20*0.25=`5Mbps`) the allotted `5Mbps` on the link between server and Router 1. Motivated by NDN consistently reporting higher avg. bitrate, we also report bandwidth utilization by NDN versus IP streaming clients as presented

in Tab. 3. If we define bandwidth utilization as the percentage of average playback bitrate per client for bottleneck bandwidth allotted per client, then NDN utilizes up to 450% of the available bandwidth (40*0.78=31.2Mbps out of 6.6Mbps) whereas IP can only utilize up to 100% (20*0.17=3.4Mbps out of 3.4Mbps) of it in the best case scenario (computed from experiment IIb results). These results indicate NDN's superior bandwidth utilization over IP across all scales. This clearly demonstrates the scalability and the benefits of using NDN for live streaming applications.

Table 2: Average QoE metric results from streaming sessions accross different experiment setup

Scenario	Rebuf %	#QS	Avg. bitrate	Spectrum	VMAF
Exp I: 4 NDN, 4 IP, 50/50 BW split, 500MB Cache					
NDN	0.2	66.9	2.74	925.1	96.9
IP	0.1	52.9	1.20	446.9	82.2
Exp IIa: 20 NDN, 20 IP, 50/50 BW split, 0MB Cache					
NDN	5.44	103.33	0.39	164.06	67.87
IP	0.71	33.6	0.25	25.96	43.74
20 NDN, 20 IP, 50/50 BW split, 250MB Cache					
NDN	3.39	152.06	0.75	710.90	66.97
IP	0.64	33.31	0.25	22.99	47.32
20 NDN, 20 IP, 50/50 BW split, 500MB Cache					
NDN	3.36	145.58	0.75	685.35	70.9
IP	0.58	34.18	0.25	24.56	47.82
Exp IIb: 40 NDN, 20 IP, 66/33 BW split, 500MB Cache					
NDN	6.24	147.03	0.78	749.28	65.03
IP	7.14	22.49	0.17	17.39	13.75

Table 3: Comparison of the theoretical available average bandwidth for each client with the average playback bitrate observed by each client.

Layer 3	Total # of clients	Bottleneck BW split	Theoretical average BW	Avg. playback bitrate	BW utilization
NDN	4NDN-4IP (Σ 8)	50%/50%	1.25	2.74	219%
IP	4NDN-4IP (Σ 8)	50%/50%	1.25	1.2	96%
NDN	20NDN-20IP (Σ 40)	50%/50%	0.25	0.75	300%
IP	20NDN-20IP (Σ 40)	50%/50%	0.25	0.25	100%
NDN	40NDN-20IP (Σ 60)	66%/33%	0.17	0.78	458%
IP	40NDN-20IP (Σ 60)	66%/33%	0.17	0.17	100%

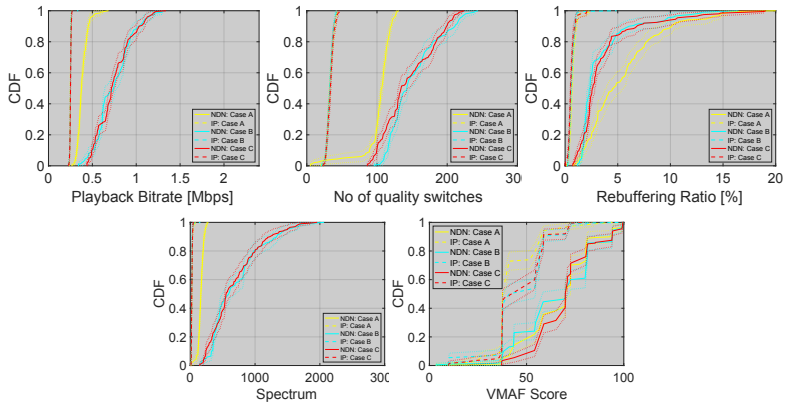

Fig. 6: Cumulative Distribution Functions for QoE metrics of 20 NDN client cases with varying cache sizes. Case A: NDN cache size 0MB; Case B: NDN cache size 250MB; Case C: NDN cache size 500MB

4.2.2. *Comparison of Perceived video quality*

Traditional metrics might not always correlate with a user's perception of a good video quality. It is imperative to test the quality of the video being streamed as well as its playback quality. The VMAF metric (as previously described in 3) predicts a viewer's perception of a video streaming quality in the form of a score between 0 to 100. Since, VMAF is reference-based, a score of 100 in our evaluation denotes video streamed at best available quality in the dataset. As we compare the quality of the video streamed at NDN and IP clients for each experiment in Tab. 2, it can be observed that NDN consistently reports higher VMAF than IP with same share of available bandwidth per client. The individual scores decrease with increase in number of clients (65 for 40 NDN clients in experiment IIb vs. 96.9 for 4 NDN clients in experiment I). The correlation of available bandwidth to VMAF is unclear, as seen with 0-cache and 250 MB cache results in IIa. Here, increased cache and ABR do not necessarily lead to increased VMAF. However it is clear that NDN's caching benefits lead to a better user viewing experience compared to IP.

Figures 5, 6, and 7 also report higher average VMAF in the case of NDN clients as compared to IP clients for experiment I, IIa & IIb respectively. In other words, NDN clients consistently streamed better perceivable quality video than IP clients.

4.2.3. *Effect of caching*

In this section, we present the effect of varying cache sizes for NDN clients in a large-scale scenario (experiment IIa). Comparing the results for the three different cache-sizes, we make three major observations. First, increasing the cache size beyond `500MB` does not lead to a further increase in QoE. Even the increase from

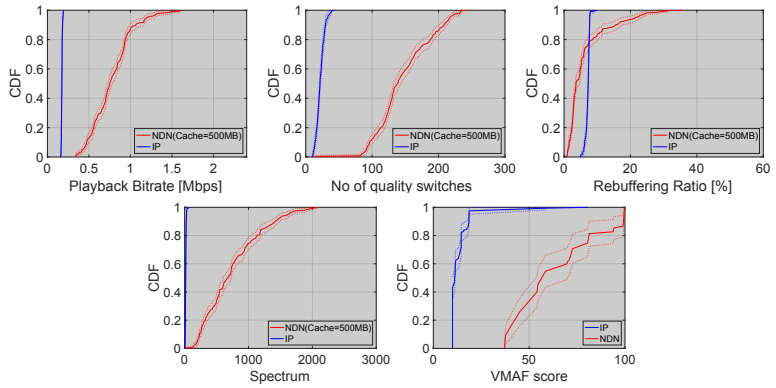

Fig. 7: Cumulative Distribution Functions for QoE metrics for the case of 40 NDN clients and 20 IP clients.

`250MB` to `500MB` results in marginal improvement of the QoE metrics. Third, the rebuffering percentage is almost identical for all three cases (see Figure 6), but slightly increasing #QS indicating interdependence between these QoE metrics. Fourth, as shown in Figure 6, the #QS metric is lowest for the 0MB cache and increases for the `250MB` and `500MB` cache cases. Compared to IP, the #QS metric is higher for all three caching scenarios and the difference to NDN is larger than in the small-scale scenario. From the large-scale scenarios in Table 2, we observe that #QS increases for NDN and decreases for IP. In the next section, we explain the reason behind this with a hypothesis as well as proposed solutions.

4.2.4. *Effect of Oscillation effect on Rebuffering, #QS & Spectrum: Causes & Solution*

Examining the additional QoE parameters shows that the increased playback bitrate and higher VMAF in the NDN case comes with the trade-off of increases in quality switches and spectrum. While the rebuffering percentage stays comparable to the IP scenario in small-scale results (experiment I), it increases in comparison with IP for large-scale cases (experiment IIa & b). This increase might have a negative impact on the viewer's QoE that can outweigh the positive impact of an increased playback bitrate. Hence, in this section we first identify the underlying cause and then suggest a solution to this problem.

Cause/Hypothesis: First, higher #QS in NDN as opposed to IP can in part be explained by examining the average bandwidth available to each IP client on the bottleneck link. In the case of 20 IP clients the average bandwidth is `0.25Mbps` per client. This bandwidth is only sufficient to stream the lowest out of 10 playback bitrates (`0.15 Mbps`). For NDN live streaming with `250MB` cache size, the average bitrate is `0.75Mbps`, which is sufficient to stream the four lowest playback bitrates. This clearly demonstrates higher opportunities for bitrate quality changes in NDN

Table 4: Avg. QoE and hit-miss ratio at respective caches for Livestreaming on NDN with 10 qualities vs. 4 qualities

Layer 3	Rebuff%	#QS	ABR	Spectrum	VMAF	Hit-Rate at Router1	Hit-Rate at Router2	Hit-Rate at Router3
NDN (10-qualities)	3.36	145.58	0.75	685.35	70.9	0.09	0.17	0.21
IP	0.58	34.18	0.25	24.56	47.8	NA	NA	NA
NDN (4-qualities)	0.1	10.5	1.47	14.3	78.6	0.12	0.63	0.72
IP	0.6	34.2	0.25	25.0	48.3	NA	NA	NA

as opposed to IP. Second, we conjecture that some of the decreases in QoE are caused by the interplay of ABR streaming and NDN. As presented by Grandl et al. [10], the potentially random placement of video segment on either the server or the caches can lead to so-called "oscillation effects". For example, if a client receives a video segment in low quality from a cache, the measured download rate might be high. Based on this observation, the ABR algorithm at the client (e.g., BOLA [43]) might decide to request the next segment in a higher quality. If this segment is currently not stored at the cache, the client has to retrieve the segment from the server and most likely experiences a lower download rate. This results in the client requesting the next segment at a (much) lower quality. This alternate retrieval of segments from server or cache can happen several times during a streaming session leading to increased #QS and spectrum in the case of NDN. Furthermore, these suspected oscillation effects caused by potential low hit rates on the caches lead to lower download rates and hence higher rebuffering percentages. This led us to further investigate the effect of available bitrates on cache hit-ratios causing potential"oscillation effects" impacting QoE. *Confirmation:* To gain more insight in the correlation between a higher number of available bitrates (in the case of NDN), the oscillation effect caused by low cache hit-rates, and the resulting QoE (#QS, spectrum & Rebuffering percentage), we first analyze the hit rate (per NDN segment) on the three routers (1-3) used in the topology for these experiments. In an additional experiment for the 20NDN/20IP client case, we set the cache sizes on all three routers to 2GB (large enough to cache all DASH segment in all playback bitrates of the video which totals to 1.4GB). We observed that the hit rates, surprisingly, do not increase with an increase in cache size and are consistently low. Reported hit-rates were 0.09, 0.22, and 0.25 for `250MB` caches, 0.09, 0.17, and 0.22 for `500MB` caches and 0.08, 0.23, and 0.19 for `2GB` caches in Router1, Router2, and Router3, respectively. The most probable reason behind low hit rates would be that all 20 NDN clients request a very disjunct set of qualities for individual DASH video segments (every segment is available in ten different bitrates (see Sect. 3.1)). Fig. 8 further proves this hypothesis by showing the number of unique qualities per segment that

are requested by the clients in an entire streaming session. With a maximum of

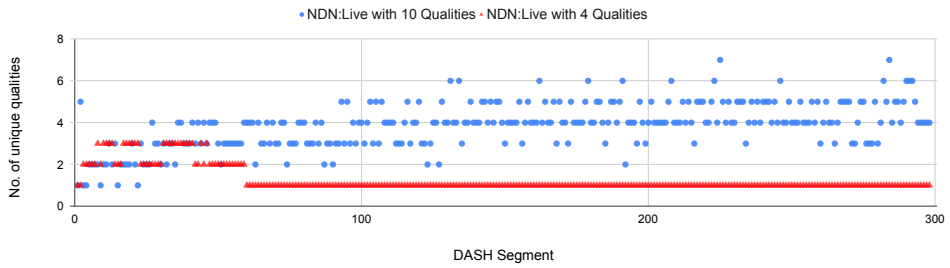

Fig. 8: Distinct qualities requested per DASH video segment by an NDN client when #qualities available at the server are 10 vs. 4. Total DASH segments in our streaming video are 298. Maximum possible unique qualities requested in both cases would be 10 and 4 respectively and minimum would be 1.

ten qualities available, NDN clients requested up to seven different qualities for a given DASH segment and an average of five different qualities for all DASH segments. This seemingly high variation in requested qualities results in low hit rates. With ten clients connected to Routers 2 and 3 each, almost every client requested a different quality per DASH segment. *Solution:* Based on these observations, we reduced the available playback bitrates from ten to four (`0.15 Mbps, 0.45Mbps, 0.92 Mbps and 1.54 Mbps`) for NDN live streaming and conducted an experiment with 20 NDN and 20 IP clients (similar to Experiment IIa in Sect. 4) and cache sizes set to `500MB` at the routers. The IP clients are still able to select from all ten playback bitrates. We observe that the number of distinct qualities were much lower in the case of NDN. Previously with ten bitrates, *mode* for the distinct qualities requested per DASH segment was four and a *maximum* of seven distinct qualities. But with four available bitrates, the *mode* reduces to one with a *maximum* of three distinct qualities requested by all NDN clients. More importantly, this resulted in higher hit rates as reported in Table 4. Figures 9 and Table 4 further confirm our hypothesis that the reduction of available qualities has a significant impact on QoE. From the CDF graphs in Figure 9, we observe that when NDN chooses from 4 qualities, it results in higher playback bitrate, lower quality switches, lower rebuffering percentage, spectrum as compared to when NDN chooses from more available qualities (in our case, 10). The VMAF score also improves with our solution indicating better visual quality of the streamed video. Furthermore, when NDN clients choose from fewer available bitrates, they outperform IP clients in terms of all QoE metrics with reduced #QS, rebuffering ratio, spectrum and increased ABR & VMAF. Additionally, we ran similar experiments for I & IIb (4NDN and 4IP, 40NDN and 20IP clients respectively) with reduced available playback bitrates around their expected average. For each case, reducing available bitrates improves all QoE metrics

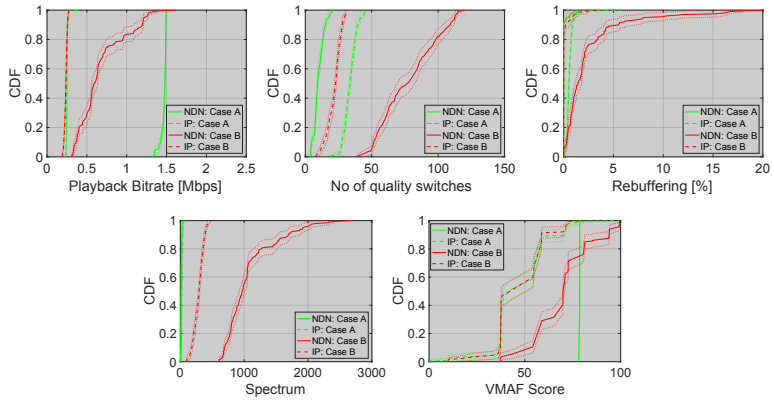

Fig. 9: Cumulative Distribution Functions for QoE metrics for the case of limited playback bitrates. Case A: NDN chooses from 4 qualities & IP from 10. Case B: Both NDN & IP choose from 10 qualities

of NDN clients as compared to IP as well as when compared to NDN clients streaming with full set of qualities. For IIb, VMAF score improved from 65 to 77, higher ABR of 1.4Mbps and lower #QS, Rebuffering % and Spectrum. These results conclusively show that with the suggested modifications, NDN outperforms IP across all QoE metrics and across different scales of experiments. Clearly, the selection of the playback bitrate for NDN live streaming was informed by the results presented in Sect. 4.2.1 and Table 2. With an average bitrate of `0.75Mbps` in the case of `500MB` cache size, selecting playback bitrates was straight-forward. To make such an approach feasible for an actual system, an approach could be implemented that collects the average client bandwidth and adapts the playback bitrates that are advertised in the MPD file once a sufficient amount of data has been collected.

4.2.5. *NDN vs. IP with client-side bottleneck*

Our motivation to enforce a server-side bottleneck so far was to observe how the architecture responds to reduced midgress traffic. Obviously, the last hop to the client can also be a bottleneck (especially in mobile scenarios). Hence, a comparative QoE analysis of NDN and IP clients streaming over a network with client-side bottleneck is also required to provide a complete picture of NDN's superior performance in the case of live-streaming. To further study such a scenario, we increased the bandwidth of the server side link (see Figure 2) to 1Gbps and added a client-side bottleneck of 10Mbps. For equal distribution of resources and a fair comparison, 5 clients were run on each node (both IP and NDN) and all nodes were connected to routers 2 & 3 with 10 Mbps links. As Figure 10 shows, NDN outperforms IP across all QoE metrics with 4 qualities. Even though the caching is limited with this client-side bottleneck, NDN still benefits from it as well as from its inherent nature of multicast and retransmissions to the nearest cache. It is also interesting to note that contrary to the results shown in Figure 9, QoE for 10 qualities is more comparable

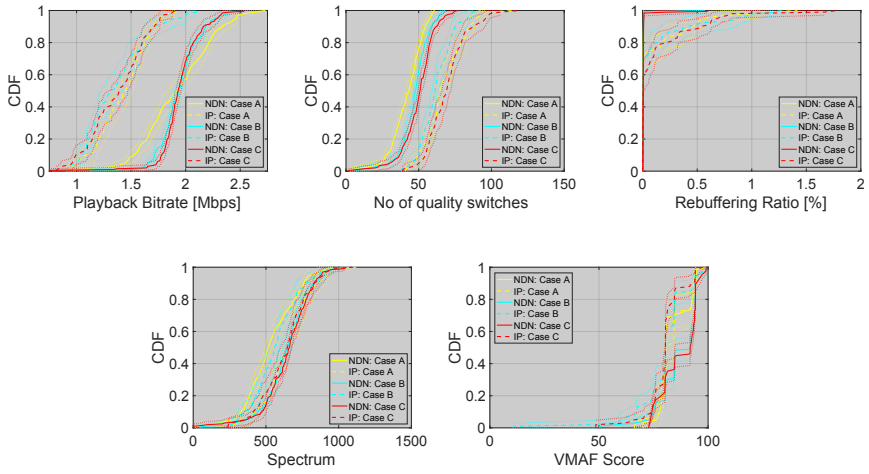

Fig. 10: Cumulative Distribution Functions for QoE metrics with client-side bottleneck. Case A: NDN chooses from 4 qualities & IP from 10. Case B: Both NDN & IP choose from 10 qualities. Case C: Case B with 1% packet loss

to the one with 4 qualities (0.15 Mbps, 0.92 Mbps, 2.89 Mbps, 3.8 Mbps) around an ideal average of 2Mbps. In this case, the bottleneck at the last hop dampens the oscillation effect. The bitrate is slightly increased in the case of 10 qualities because of more available qualities which in turn slightly worsens #QS, spectrum, and rebuffering percentage. In the event of loss, it is intuitive that the QoE will be lower than in the lossless cases. However even without reduced bitrates, Figure 10 shows that NDN still outperforms IP with respect to all QoE metrics besides the spectrum due to increased magnitude of variability with 10 qualities. NDN clients also consistently report higher VMAF across all cases than IP.

5. Related Work

5.1. *CDNs and NDN*

To accomplish large-scale live streaming, CDN providers need to *work around* several limitations of the TCP/IP architecture such as lack of native IP multicast support, lack of support for caching at the network layer, and more. First, they need to create a data delivery infrastructure that can scale to the enormous demand. Second, TCP/IP does not support easy and efficient multicast. Finally, TCP/IP does not support caching of content at the network layer. Caching at the application layer (e.g., HTTP caches) also introduces several issues. For example, clients must find and connect to the appropriate caches, the CDNs must strategically place content in the caches, and continuously monitor for failure or service degradation. NDN, on the other hand, provides several desirable properties for live streaming.

First, NDN provides native multicast at the network layer through Interest aggregation and caching, caches Data packets at the network layer, and does not require complex infrastructure setup beforehand. In the case of cache or link failure, NDN is able to mitigate the failure by forwarding traffic around it (if there are multiple paths) without application or operator intervention. These properties can be utilized for live-streaming - both inside a CDN infrastructure and user-based live streaming. A CDN can certainly deploy several NDN based live-streaming servers inside its infrastructure that will work in parallel with existing IP-based streaming mechanisms. Previous work by Ghasemi et al. [9] compared an NDN based video streaming solution deployed on the NDN Testbed with two well-known CDNs, Akamai and Fastly. While this work did not attempt a hybrid solution, it showed NDN can reduce origin workload compared to traditional CDNs. Another work by Thelagathoti et al. [46] demonstrated that NDN deployment inside a CDN infrastructure will help with efficient data retrieval and improve QoE.

5.2. *SDN, NFV, and NDN*

Several works have highlighted the benefits of an architecture that integrates SDN with NFV [6, 24]. For instance, the agility this integration provides to infrastructure and network service design can be very desirable for any dynamic and scalable architectural framework [30]. NFV and SDN combined have revolutionized network architectures that are able to cope with the continuous growth in data-traffic [28]. They provide the ability to virtualize any network infrastructure based on its requirements. Hence, the decision to use our SDN-NFV setup. There has been work on SDN-NFV infrastructures that handle heterogeneous network technologies (e.g., [22, 47, 48]) but none of them compare multiple network stacks and their performance. In [21], Mai et al. have implemented NDN technologies with SDN-NFV support but the novelty of our work lies in the heterogeneity of the network protocol stacks as implemented in [44]. Performance analysis over IP versus non-IP protocols [17] or specifically IP versus NDN protocols [38, 36] have been executed before but not with the design flexibility that comes with the benefits of SDN programmability [44]. Kanada et al. [17] use virtual link tunnels to encapsulate IP and non-IP packets whereas Satria et al. [38] evaluate them separately and not in the context of video streaming in a non-virtualized setup.

Several works have proposed translation between TCP/IP and NDN so that they can coexist. Moiseenko et al. proposed TCP over ICN where TCP traffic is converted into ICN traffic [27]. Shannigrahi et al. proposed IPoC [40] where TCP/IP traffic is encapsulated into NDN packets for transport. Refaei et al. proposed an IP-ICN gateway that allows IP client-server communication to operate seamlessly through an NDN cloud [34]. Nour et al. [29] propose an approach that uses NFV for ICN/IP hybrid routers that require predefining a set of regions. On the contrary, we utilize layer 2 header information to indicate different types of traffic that allows us to differentiate between IP and NDN traffic in real-time and decide whether to forward it to IP or NDN data sourceS.

6. Conclusion & Future Work

In this paper, we propose and show the benefits of a hybrid and flexible streaming approach which supports multiple internet architectures over a traditional IP approach for improved QoE in live adaptive bitrate video streaming applications. To achieve this goal our approach employs network and containerization techniques. We present a detailed description to demonstrate how SDN, NFV, kernel virtualization and containerization can be orchestrated to provide a hybrid and highly scalable streaming architecture. We implement this setup in the CloudLab testbed and perform an extensive performance evaluation of our hybrid streaming approach where we not only meausre the playback quality but also the video quality itself. The evaluation results demonstrate the profit in terms of average bitrate, bandwidth utilization and better video quality(VMAF) from our approach and reveal that live streaming can be performed efficiently, is scalable, and provides good QoE with the help of NDN. Using containers for the NDN streaming clients provides a method that can activate such clients without end user involvement. Counterintuitive to experience gained in the case of ABR streaming over TCP/IP, we show that a reduced set of available playback bitrates leads to better performance in the case of NDN-based live streaming and outperforms IP-based live streaming under both server and client-bottleneck conditions and across different scale of clients. We also show that NDN-based live streaming behaves fairly to IP-based session and does not negatively impact the QoE of these sessions. In future work we plan to to develop new ABR algorithms that are cognizant that NDN provides in-network caching. We will also study if the approach of caching at the level of NDN Segments is appropriate in the case of ABR streaming.

Acknowledgements

This work was funded by National Science Foundation (NSF) grants CNS 19-01137, OAC-2019163, OAC-2126148, and OAC-2019012. All opinions and statements in the above publication are of the authors and do not represent NSF positions.

References

[1] S. H. Ahmed, S. H. Bouk, and D. Kim. *Content-Centric Networks: An Overview, Applications and Research Challenges*. Springer Publishing Company, Incorporated, 1st edition, 2016.

[2] T. Anderson, K. Birman, R. Broberg, M. Caesar, D. Comer, C. Cotton, M. J. Freedman, A. Haeberlen, Z. G. Ives, A. Krishnamurthy, et al. The nebula future internet architecture. In *The Future Internet Assembly*, pages 16–26. Springer, 2013.

[3] Docker. Docker. `"https://www.docker.com/"`, 2020.

[4] Docker. Host networking. `"https://docs.docker.com/network/host/"`, 2020.

[5] Docker. Macvlan networks. `"https://docs.docker.com/network/macvlan/"`, 2020.

[6] Q. Duan, N. Ansari, and M. Toy. Software-defined network virtualization: an architectural framework for integrating sdn and nfv for service provisioning in future networks. *IEEE Network*, 30(5):10–16, 2016.

[7] D. Duplyakin, R. Ricci, A. Maricq, G. Wong, J. Duerig, E. Eide, L. Stoller, M. Hibler, D. Johnson, K. Webb, A. Akella, K. Wang, G. Ricart, L. Landweber, C. Elliott,

M. Zink, E. Cecchet, S. Kar, and P. Mishra. The design and operation of cloudlab. In *2019 USENIX Annual Technical Conference (USENIX ATC 19)*, pages 1–14, Renton, WA, July 2019. USENIX Association.

[8] L. Foundation. tc. "https://linux.die.net/man/8/tc", 2020.

[9] C. Ghasemi, H. Yousefi, and B. Zhang. Far cry: Will cdns hear ndn's call? In *Proceedings of the 7th ACM Conference on Information-Centric Networking*, ICN '20, page 89–98, New York, NY, USA, 2020. Association for Computing Machinery.

[10] R. Grandl, K. Su, and C. Westphal. On the interaction of adaptive video streaming with content-centric networking. In *2013 20th International Packet Video Workshop*, pages 1–8, 2013.

[11] P. Gusev, Z. Wang, J. Burke, L. Zhang, T. Yoneda, R. Ohnishi, and E. Muramoto. Real-time streaming data delivery over named data networking. *IEICE Transactions on Communications*, 99(5):974–991, 2016.

[12] S. Ha, I. Rhee, and L. Xu. Cubic: A new tcp-friendly high-speed tcp variant. *SIGOPS Oper. Syst. Rev.*, 42(5):64–74, July 2008.

[13] N. Handigol, B. Heller, V. Jeyakumar, B. Lantz, and N. McKeown. Reproducible network experiments using container-based emulation. In *Proceedings of the 8th International Conference on Emerging Networking Experiments and Technologies*, CoNEXT '12, page 253–264, New York, NY, USA, 2012. Association for Computing Machinery.

[14] A. Inc. Apple HTTP Live Streaming. "https://developer.apple.com/streaming/", 2020.

[15] A. S. Incorporated. Adobe HTTP Dynamic Streaming. "http://www.adobe.com/products/hds-dynamic-streaming.html", 2020.

[16] P. Juluri. Astream. "https://github.com/pari685/AStrea", 2020.

[17] Y. Kanada and A. Nakao. Development of a scalable non-ip/non-ethernet protocol with learning-based forwarding method. In *2012 World Telecommunications Congress*, pages 1–6, 2012.

[18] Z. Kong, X. Ma, Y. Zhang, Z. Zhang, D. Pesavento, S. Shannigrahi, S. Dulal, and J. Shi. ndn-python-repo. "https://github.com/UCLA-IRL/ndn-python-repo", 2020.

[19] LBL. Singularity. "https://sylabs.io/docs/", 2020.

[20] Z. Li, A. Aaron, I. Katsavounidis, A. Moorthy, and M. Manohara. Toward a practical perceptual video quality metric. "https://netflixtechblog.com/toward-a-practical-perceptual-video-quality-metric-653f208b9652", 2016.

[21] H. L. Mai, M. Aouadj, G. Doyen, W. Mallouli, E. M. de Oca, and O. Festor. Toward content-oriented orchestration: Sdn and nfv as enabling technologies for ndn. In *2019 IFIP/IEEE Symposium on Integrated Network and Service Management (IM)*, pages 594–598, 2019.

[22] R. Martínez, A. Mayoral, R. Vilalta, R. Casellas, R. Muñoz, S. Pachnicke, T. Szyrkowiec, and A. Autenrieth. Integrated sdn/nfv orchestration for the dynamic deployment of mobile virtual backhaul networks over a multilayer (packet/optical) aggregation infrastructure. *IEEE/OSA Journal of Optical Communications and Networking*, 9(2):A135–A142, 2017.

[23] S. Mastorakis, A. Afanasyev, and L. Zhang. On the evolution of ndnsim: An open-source simulator for ndn experimentation. *SIGCOMM Comput. Commun. Rev.*, 47(3):19–33, Sept. 2017.

[24] J. Matias, J. Garay, N. Toledo, J. Unzilla, and E. Jacob. Toward an sdn-enabled nfv architecture. *IEEE Communications Magazine*, 53(4):187–193, 2015.

[25] D. Merkel. Docker: lightweight linux containers for consistent development and deployment. *Linux journal*, 2014(239):2, 2014.

[26] Microsoft. Microsoft Silverlight Smooth Streaming. "https://www.microsoft.com/

silverlight/smoothstreaming", 2020.

[27] I. Moiseenko and D. Oran. Tcp/icn: Carrying tcp over content centric and named data networks. In *Proceedings of the 3rd ACM Conference on Information-Centric Networking*, ACM-ICN '16, page 112–121, New York, NY, USA, 2016. Association for Computing Machinery.

[28] V. Nguyen, A. Brunstrom, K. Grinnemo, and J. Taheri. Sdn/nfv-based mobile packet core network architectures: A survey. *IEEE Communications Surveys Tutorials*, 19(3):1567–1602, 2017.

[29] B. Nour, F. Li, H. Khelifi, H. Moungla, and A. Ksentini. Coexistence of icn and ip networks: an nfv as a service approach. In *2019 IEEE Global Communications Conference (GLOBECOM)*, pages 1–6. IEEE, 2019.

[30] N. Omnes, M. Bouillon, G. Fromentoux, and O. L. Grand. A programmable and virtualized network it infrastructure for the internet of things: How can nfv sdn help for facing the upcoming challenges. In *2015 18th International Conference on Intelligence in Next Generation Networks*, pages 64–69, 2015.

[31] R. Pauly, C. Ogle, C. Mcknight, D. Reddick, J. Presley, S. Shannigrahi, and A. Feltus. NDN-TR68: Utilizing NDN for Domain Science Applications - a Genomics Example - Named Data Networking (NDN), Mar 2021. [Online; accessed 8. Mar. 2021].

[32] B. Pfaff, J. Pettit, T. Koponen, E. J. Jackson, A. Zhou, J. Rajahalme, J. Gross, A. Wang, J. Stringer, P. Shelar, K. Amidon, and M. Casado. The design and implementation of open vswitch. In *Proceedings of the 12th USENIX Conference on Networked Systems Design and Implementation*, NSDI'15, page 117–130, USA, 2015. USENIX Association.

[33] J. J. Quinlan, A. H. Zahran, and C. J. Sreenan. Datasets for avc (h.264) and hevc (h.265) evaluation of dynamic adaptive streaming over http (dash). In *Proceedings of the 7th International Conference on Multimedia Systems*, MMSys '16, pages 2–3, New York, NY, USA, 2016. Association for Computing Machinery.

[34] T. Refaei, J. Ma, S. Ha, and S. Liu. Integrating ip and ndn through an extensible ip-ndn gateway. In *Proceedings of the 4th ACM conference on information-centric networking*, pages 224–225, 2017.

[35] Y. Ren, J. Li, S. Shi, L. Li, G. Wang, and B. Zhang. Congestion control in named data networking – a survey. *Computer Communications*, 86:1–11, 2016.

[36] J. Samain, G. Carofiglio, L. Muscariello, M. Papalini, M. Sardara, M. Tortelli, and D. Rossi. Dynamic adaptive video streaming: Towards a systematic comparison of icn and tcp/ip. *IEEE Transactions on Multimedia*, 19(10):2166–2181, 2017.

[37] J. Samain, G. Carofiglio, L. Muscariello, M. Papalini, M. Sardara, M. Tortelli, and D. Rossi. Dynamic adaptive video streaming: Towards a systematic comparison of icn and tcp/ip. *IEEE Transactions on Multimedia*, 19(10):2166–2181, 2017.

[38] M. N. D. Satria, F. H. Ilma, and N. R. Syambas. Performance comparison of named data networking and ip-based networking in palapa ring network. In *2017 3rd International Conference on Wireless and Telematics (ICWT)*, pages 43–48, 2017.

[39] I. Seskar, K. Nagaraja, S. Nelson, and D. Raychaudhuri. Mobilityfirst future internet architecture project. In *Proceedings of the 7th Asian Internet Engineering Conference*, pages 1–3, 2011.

[40] S. Shannigrahi, C. Fan, and G. White. Bridging the icn deployment gap with ipoc: An ip-over-icn protocol for 5g networks. In *Proceedings of the 2018 Workshop on Networking for Emerging Applications and Technologies*, pages 1–7, 2018.

[41] I. Sodagar. The MPEG-DASH standard for multimedia streaming over the Internet. *IEEE MultiMedia*, 18(4):62–67, April 2011.

[42] K. Spiteri, R. Sitaraman, and D. Sparacio. From theory to practice: Improving bitrate

adaptation in the dash reference player. *ACM Trans. Multimedia Comput. Commun. Appl.*, 15(2s):5–6, July 2019.

[43] K. Spiteri, R. Urgaonkar, and R. K. Sitaraman. Bola: Near-optimal bitrate adaptation for online videos. In *IEEE INFOCOM 2016 - The 35th Annual IEEE International Conference on Computer Communications*, pages 1–9, 2016.

[44] B. Suresh, D. Bhat, and M. Zink. An evaluation of sdn and nfv support for parallel, alternative protocol stack operations. In *2018 IEEE International Conference on Communications (ICC)*, pages 1–7, 2018.

[45] I. Technical Report. Cisco Systems. Cisco. 2020. cisco visual networking index (vni) complete forecast update. 2017 - 2022.

[46] R. K. Thelagathoti, S. Mastorakis, A. Shah, H. Bedi, and S. Shannigrahi. Named data networking for content delivery network workflows. In *2020 IEEE 9th International Conference on Cloud Networking (CloudNet)*, pages 1–7. 10.1109/CloudNet51028.2020.9335806, 2020.

[47] R. Vilalta, A. Mayoral, R. Casellas, R. Martínez, and R. Muñoz. Experimental demonstration of distributed multi-tenant cloud/fog and heterogeneous sdn/nfv orchestration for 5g services. In *2016 European Conference on Networks and Communications (EuCNC)*, pages 52–56, 2016.

[48] R. Vilalta, A. Mayoral, R. Muñoz, R. Casellas, and R. Martínez. Multitenant transport networks with sdn/nfv. *Journal of Lightwave Technology*, 34(6):1509–1515, 2016.

[49] L. Zhang, A. Afanasyev, J. Burke, V. Jacobson, k. claffy, P. Crowley, C. Papadopoulos, L. Wang, and B. Zhang. Named data networking. *SIGCOMM Comput. Commun. Rev.*, 44(3):66–73, July 2014.

[50] M. Zink, J. Schmitt, and R. Steinmetz. Layer-encoded video in scalable adaptive streaming. *IEEE Transactions on Multimedia*, 7(1):75–84, 2005.