

Central Lancashire Online Knowledge (CLoK)

Title	Improving video QoE with IP over ICN
Type	Article
URL	https://clock.uclan.ac.uk/30425/
DOI	https://doi.org/10.1002/nem.2057
Date	2018
Citation	Doumanis, Ioannis, Phinikarides, Alexander, Xylomenos, George, Porter, Stuart C.M. and Georgiades, Michael (2018) Improving video QoE with IP over ICN. International Journal of Network Management. e2057. ISSN 1055-7148
Creators	Doumanis, Ioannis, Phinikarides, Alexander, Xylomenos, George, Porter, Stuart C.M. and Georgiades, Michael

It is advisable to refer to the publisher's version if you intend to cite from the work.
<https://doi.org/10.1002/nem.2057>

For information about Research at UCLan please go to <http://www.uclan.ac.uk/research/>

All outputs in CLoK are protected by Intellectual Property Rights law, including Copyright law. Copyright, IPR and Moral Rights for the works on this site are retained by the individual authors and/or other copyright owners. Terms and conditions for use of this material are defined in the <http://clock.uclan.ac.uk/policies/>

RESEARCH ARTICLE

Improving Video QoE with IP-over-ICN

Ioannis Doumanis¹ | Alexander Phinikarides² | George Xylomenos³ | Stuart CM Porter⁴ | Michael Georgiades²

¹University of Central Lancashire, United Kingdom

²PrimeTel PLC, Cyprus

³Athens University of Economics and Business, Greece

⁴Consultant, United Kingdom

Correspondence

*George Xylomenos, Athens University of Economics and Business, Patission 76, Athens 104 34, Greece. Email: xgeorge@aueb.gr

Abstract

Information-centric networking (ICN) has long been advocating for radical changes to the Internet, but the upgrade challenges that this entails have hindered its adoption. To break this loop, the POINT project proposed a hybrid, IP-over-ICN, architecture: IP networks are preserved at the edge, connected to each other over an ICN core. This exploits the key benefits of ICN, enabling individual network operators to improve the performance of their IP-based services, without changing the rest of the Internet. This paper first provides an overview of POINT and outlines how it can improve upon IP in terms of performance and resilience. It then describes a trial of the POINT prototype in a production network, where real users operated actual IP-based applications. As part of the trial, we carried out experiments to evaluate the Quality of Experience (QoE) for video services offered via either HLS or IPTV, using either IP or POINT as a substrate. The results from the trial verify that the IP-over-ICN approach of POINT offers enhanced QoE to the users of these video services, compared to traditional IP, especially under exceptional network conditions.

KEYWORDS:

ICN, POINT, HLS, IPTV, QoE, Trials

1 | INTRODUCTION

Information-Centric Networking (ICN)¹ proposes replacing the endpoint-based communication of the current Internet with an architecture focusing on the exchange of named data. An active research community has grown around ICN^{2,3}, following different design approaches, but as ICN requires overhauling the entire Internet infrastructure *and* rewriting all applications, it is hard to make the transition from research testbeds to operational networks. Based on their experience in pioneering ICN projects, including PSIRP, PURSUIT and COMET, the partners of the EU-funded H2020 POINT project¹ have long realized the difficulty of replacing a hugely successful incumbent architecture with a promising, but unproven, clean-slate one. Therefore, POINT has taken a more pragmatic approach to the introduction of ICN: instead of replacing the Internet wholesale, POINT proposes supporting the existing IP-based applications and services over an ICN core network⁴. POINT thus addresses the needs of individual operators who want to retain compatibility with the existing Internet, while taking advantage of specific ICN solutions to improve the performance of their network.

Video services, which contribute the majority of Internet traffic⁵, were the primary target for improvement in POINT; specifically, we considered IPTV-based services, commonly used to deliver live TV, and HLS-based services, commonly used for

¹<https://www.point-h2020.eu/>

on-demand video streaming. In the case of IPTV, its efficient support requires operators to use non-standard extensions to support IP multicast, such as IGMP snooping, effectively limiting such services to single operator networks. With POINT and its native multicast capability, IPTV applications can be supported with plain SDN switches, while gaining flexibility in terms of traffic handling. In the case of HLS, video is delivered individually to each user over HTTP, even when many users request the exact same pieces of content. With POINT, multiple requests for the same content can be aggregated transparently to the endpoints and transported via native multicast, while surrogate servers can be quickly added and dropped for load balancing or fault tolerance purposes, via the tight integration of POINT with SDN.

To showcase the benefits of POINT and its compatibility with existing IP-based networks and services, POINT executed a trial on PrimeTel's network in Cyprus, with real users accessing both IPTV and HLS-based video services over either IP or POINT⁶, under regular and exceptional conditions. In addition to monitoring the *Quality of Service* (QoS) offered by each type of network, we designed an extensive set of experiments to assess the *Quality of Experience* (QoE) of the video services, as perceived by the trial participants. This paper expands our previous reports about that trial^{6,7}, explaining in detail how the experiments were designed and carried out, as well as analyzing their results, which strongly support the claimed benefits of POINT.

The outline of the remainder of this paper is as follows. In Section 2 we provide an overview of the POINT approach, including its architecture and goals. In Section 3 we describe how IPTV and HLS are implemented in the current Internet, and how they can be enhanced with POINT. We then describe the design of the POINT trial in Section 4, focusing on its QoE measurement aspects. Section 5 presents the results of the trial, in terms of both QoE and QoS. Finally, Section 6 summarizes our findings.

2 | THE POINT PROJECT

2.1 | The POINT architecture

The IP-over-ICN architecture proposed by POINT aims to replace the network of individual network operators, so as to improve their IP-based services. POINT is a drop-in network replacement: it does not require any changes to existing IP clients, servers and services, or to the IP routers/gateways of interconnected operators. This is achieved by combining an ICN core network with a set of *Network Attachment Points* (NAPs), which serve as gateways between the IP and ICN worlds. The baseline POINT architecture was derived from the PURSUIT ICN architecture², in which the endpoints can publish and subscribe to named information items. This publish/subscribe architecture is facilitated by three core functions: a *Rendezvous* (RV) function that matches publisher and subscriber nodes; a *Topology Management* (TM) function that calculates paths between the various nodes and encodes them into *Forwarding Identifiers* (FIDs); and, a *Forwarding Node* (FN) function that allows data items to be forwarded in the network based on the FIDs.

To construct the FIDs, we assign a *Link Identifier* (LID) to each network link. The LID is a long string of bits with a few bits set to 1. The FID representing the set of links that a packet must traverse, whether these form a unicast path or a multicast tree, is the bitwise OR of the corresponding LIDs. The FIDs are included in packet headers, allowing FNs to forward packets by performing a bitwise AND and a comparison between an FID and each of the LIDs of their attached links: if the FID contains all the bits of an LID, then the packet is forwarded via the corresponding link. POINT, therefore, enables multicast and anycast without any transient state in the switches, which only need to know the LIDs of their links. For small networks (as in the POINT trial), each link has a different bit set to 1 in its LID. For larger networks we use Bloom filters for the LIDs and FIDs⁸.

Figure 1 outlines the main components of the POINT architecture. The RV and TM functionalities are the main control functions of the ICN cloud. Standard *Software Defined Networking* (SDN) switches are used for the FN functionality in POINT, replacing the dedicated forwarding nodes of PURSUIT⁹. The SDN switches are unaware of POINT, therefore they are controlled by an SDN controller¹⁰, which communicates with the TM function: the TM instructs the SDN controller how to configure the SDN switches so as to translate the FIDs included in packets to forwarding actions on their attached links, while the SDN controller informs the TM of any changes in the topology and operation of the network. The exact layout of the ICN cloud is unimportant, as long as it is connected and the NAPs are at its edges.

Another new feature introduced by POINT is the fast formation of multicast trees, which is particularly important for video applications. In PURSUIT, adding or dropping a multicast subscriber from a group required communication with the RV and TM functions. POINT takes advantage of the fact that the forwarding scheme allows merging unicast paths into multicast trees by using a simple bitwise OR over the FIDs of the unicast paths; multicast senders cache unicast paths to receivers and combine them dynamically into multicast trees as needed. Whenever the SDN controller is notified by the switches that links have been




FIGURE 1 The POINT architecture.

added or dropped, it informs the TM, which uses multicast to notify the NAPs about the link changes. Each NAP checks whether these links are part of any cached paths, using the bitwise operations mentioned above; if so, the paths are invalidated.

To preserve the IP interface towards UEs and other operators, POINT uses a gateway approach, as mentioned above. The NAPs, which are the access gateways of customers to the network, or the border gateways to peering networks, handle all the protocols offered at the IP interface, either directly at the IP layer, or, if possible, at the transport or application layer, for example, HTTP. As a result, the POINT network looks like a standard IP network to UEs and peering networks.

2.2 | The goals of POINT

To verify whether POINT does indeed offer a better service than an IP core, thus offering an incentive for adoption to network operators, we must first define what *better* means⁴. While the project identified and measured a large set of *Key Performance Indicators* (KPIs), we will focus here on the performance and resiliency benefits for the video delivery services which were showcased in the POINT trial and were assessed from the QoE perspective. Current IP networks offer video delivery services using two different solutions. On the one hand, IPTV services offered by operators over their own access networks exploit IP multicasting to offer synchronized content viewing, relying on (statefull) snooping IGMP switches to support the economical distribution of many channels with fast channel switching. On the other hand, over-the-top services (e.g., YouTube) use adaptive HTTP streaming (e.g., *HTTP Live Streaming* or HLS)¹¹ to cross diverse networks. Operators also use HLS to extend their IPTV services to customers attached to other networks, for example, 3G/4G cellular networks. POINT offers multiple improvements:

- In the *performance improvement* front, with IP users must separately stream the same content over HLS, even for live events with many simultaneous viewers. POINT can transparently combine independent HLS sessions of users asking for the same content at roughly the same time, merging their individual requests as they enter the POINT network, and serving them with a single response delivered via multicast, which is turned back into individual unicasts when exiting the POINT network. This *coincidental multicast* approach, where multicast groups are spontaneously formed, provides a dramatic reduction in network load, making solutions such as HLS competitive with standard IPTV solutions for live streaming, as the overhead and redundancy of individual video streams vanishes. In principle, this could allow an operator to use HLS for both over-the-top services and native IPTV services. An evaluation of the effectiveness of coincidental multicast indicates that even when moderate numbers of users are synchronized (50–100, which is reasonable for live streaming), traffic inside the POINT network can be reduced by up to 90%, for a wide range of content popularity distributions¹². While operators can mimic HLS multicast by inserting HTTP caches in their network, achieving similar traffic reductions requires adding caches in each and every node in the network, while the POINT solution only uses plain SDN switches.
- In the *service resilience* front, the reliance of POINT on centralized path management allows network operators to quickly adapt to changing network conditions. Paths can be recalculated at a single spot (the TM) and they only need to be communicated to the edges of the POINT network to modify routing. In contrast, in IP networks routing is distributed, so changes must be slowly propagated everywhere before routing converges. POINT exploits these features to support server

and path failover, transparently to both clients and servers. For example, a scenario tested during the POINT trial switched HLS users to a different server, without disturbing the service; the equivalent task in IP requires DNS-level changes which are slow enough to visibly disrupt HLS services. POINT also developed a solution for supporting IPTV services, with bandwidth requirements and latency performance similar to those of an IP-based solution. However, POINT only relies on stateless SDN switches rather than the statefull IGMP switches used with IP. The combination of stateless switches, centralized path management and source routing in POINT allows near instantaneous rerouting of multicast trees when paths fail. A scenario tested during the POINT trial switches all traffic from a failed link to an alternative live link, without disturbing the IPTV service; the equivalent task in IP requires the switches to first rebuild their spanning tree and then rebuild their IGMP snooping state, visibly disrupting the IPTV service.

Improvements are possible even without explicit support from POINT, as all kinds of IP-based traffic can benefit from the flexible traffic engineering made possible by POINT's centralized topology management and stateless switching. However, applications where content is identified by name can be specifically optimized for POINT, for example, applications relying on the *Constrained Application Protocol* (CoAP), for which POINT offers a dedicated solution that also exploits multicast when applicable¹³. Of course, migrating to POINT requires far more than a proof of concept, therefore members of the POINT team have led an initiative to define paths for possible migration of the Internet to ICN, which includes the IP-over-ICN approach¹⁴.

3 | VIDEO SERVICES

Many ISPs offer IPTV services, consisting of a set of live TV channels delivered via the *Real-time Transport Protocol* (RTP) over UDP and IP multicast. The transmitter addresses each TV channel to a separate IP multicast address, and receivers subscribe to the multicast group corresponding to the TV channel they want to watch. As IPTV services rely on partially non-standard hardware and software, we will base our description on the IPTV service offered by PrimeTel and also used for the POINT trial. HLS services, whether for video on demand or live TV, are more standardized, since they simply use HTTP in a specific manner, but we will also base our description on the HLS service offered by PrimeTel and used for the POINT trial.




FIGURE 2 PrimeTel's IPTV logical network topology.

3.1 | IPTV

In the PrimeTel IPTV service, whose logical layout is depicted in Figure 2, the video/audio streams (TV channels) produced by the IPTV server are sent to the IPTV Core Network using a separate IP Multicast address per stream. The IPTV Core Network, an L3 (routed) network, uses *Protocol Independent Multicast* (PIM) to route these streams between core routers. The video streams exit the IPTV Core Network and enter the Metro Network, an L2 (switched) Network, which, in turn, feeds the (also L2) Access Networks (DSLAM, DSL Modem and *Set Top Box*, or STB). The Metro and Access networks use the *Internet Group Management Protocol* (IGMP)¹⁵ to signal the groups that the network should deliver to each user. The STBs receive commands from the user's remote control and issue IGMP "join" and "leave" messages to indicate which channel the user wants to receive: a channel switch translates to an IGMP "leave" for the previous channel and an IGMP "join" for the new channel. The video is decoded by the STB and shown on a standard screen.

To avoid forwarding all groups to all customers, the L2 switches at the Metro and Access Networks use *IGMP snooping*¹⁶ to read the L3 headers of the packets and detect IGMP messages. Normally, the switch would flood multicast packets to every port, which would result in broadcasting each group to all end hosts. With IGMP snooping, the switch "listens" for the exchanges of

IGMP messages and builds a list of all the ports that have requested a particular multicast group. When an IGMP “join” message arrives for a stream over port n , the stream starts being forwarded over the port; when an IGMP “leave” message arrives for that stream over port n , the stream stops being forwarded over that port. Due to IGMP snooping, only one stream per channel reaches each L2 switch, therefore the switch is responsible for copying and forwarding the stream to the ports needed. As a result, for each IPTV stream, a multicast tree is formed with an IPTV Core router as the root and the STBs as the leaves.

The POINT network also supports native multicast, but with stateless SDN switching⁹. Multicast trees are formed at the NAPs attached to the stream sources, using the centralized TM to create paths and the fast multicast formation scheme to quickly add and drop recipients. Since source routing is used, there is no state to maintain and communicate to the SDN switches. While this approach is no faster or more economical than what an IPTV service already offers, in practical networks where there are many redundant links for fault tolerance, a stateless approach is very useful when failures occur. In a regular IPTV network, a spanning tree is created over the topology and some links are disabled for traffic forwarding purposes to avoid loops. When a link fails, the switches attached to it trigger a recalculation of the spanning tree. Due to the topological change however, the state of the IGMP snooping switches is now partially invalid. As a result, the IGMP state needs to be re-established by snooping inside new IGMP messages, a time-consuming process. In contrast, in the POINT network, there is no need to recalculate the spanning tree and re-establish switch state. If one link goes down, only the NAP attached to the IPTV server must modify its routes to switch to the other.

In the trial, the POINT network replaced the Metro and Access Networks of PrimeTel, an L2 network of IGMP snooping switches. Although extending POINT to replace the entire multicast routing substrate is straightforward, the L2 network is the most interesting part, as it relies on non-standard extensions (IGMP snooping) and a specific topology (tree) to efficiently support multicast. For the POINT platform to support the traffic between the unmodified IPTV Server and STBs, the NAPs must translate IP into ICN semantics and vice versa, handling the IP Multicast packets sent from the IPTV Server to the STBs and the IGMP messages exchanged between them. Essentially, the POINT network acts as a “big” L2 IGMP Snooping switch, forwarding IPTV streams from the server-side NAP (sNAP) leading to the IPTV server, to the client-side NAPs (cNAPs) leading to the STBs.

Our solution realizes the IGMP operations for joining and leaving an IP multicast group via publications to an ICN name, representing a control channel, while sending data to the group is realized via publications to another ICN name, representing a data channel⁷. A sender to an IP multicast group acts as a subscriber to the corresponding control channel, in order to be notified about join and leave messages to that group. On the other hand, the receivers of the IP multicast group act as subscribers to the corresponding data channel, in order to receive data packets. When a join or leave message is sent by a receiver, it is translated to a publication to the control channel with an implicit subscription or unsubscription to the data channel. With implicit subscriptions and unsubscriptions, there is no need to involve the RV and TM in each operation: the sNAP simply ORs the FIDs of the unicast paths to each cNAP *currently* participating in the group to form a multicast tree. There is no need to maintain per group state inside the network, as all forwarding is based on FIDs selected at the edges, exactly as in unicast.




FIGURE 3 PrimeTel's HLS logical network topology.

3.2 | HLS

In the PrimeTel HLS service, whose logical layout is depicted in Figure 3, the video/audio streams (TV channels) produced by the HLS server are sent to the regular IP Core Network and are routed as plain IP-based traffic, using OSPF. The video streams exit the IP Core Network and enter the Metro and Access Networks, eventually reaching the DSL modem which delivers them to a computer. PrimeTel uses HLS to deliver live video on top of the Internet when IP multicast cannot be used, which is the case when users are not connected to its own distribution network. This is a common choice for operators serving customers connected to different networks, as well as for over-the-top providers, since HLS only requires HTTP, which is available universally. On

the server side, an HTTP server is combined with a streamer/encoder to segment the video and encode the chunks into a range of qualities. On the client side, a (typically Javascript-based) application decides on the quality of the video to fetch, adapting to network conditions. Since each user is served over an individual HTTP session, a farm of HLS servers is needed to serve many customers. When live video is transmitted, segmentation and encoding must take place in real time, thus the HLS server needs to rewrite periodically the playlist files which define the available video qualities and the list of available video chunks.

When HLS traffic crosses the POINT network, the POINT NAPs reduce the amount of traffic by exploiting coincidental multicast for HTTP messages. Since HLS clients choose the quality of the next video chunk depending on their individually perceived network conditions, there is no way to predict what they will ask for next. Instead, multicast groups are formed dynamically as requests arrive at the sNAP, exploiting POINT's capability for very fast multicast group formation. An HTTP request is translated by the cNAP to a publication to a name corresponding to the server's DNS name, with an implicit subscription to a name corresponding to the parameters of the request. As a result, when multiple HTTP requests for the same content piece are received, their response is automatically sent via multicast by the sNAP to all the cNAPs waiting for it. The HTTP handler at the sNAP adds a slight delay when responding to HLS requests, to increase the chance of finding more similar requests to serve as part of a multicast transmission¹⁷. The HTTP responses are then converted back to regular unicast HTTP messages by the cNAPs.

In addition to saving bandwidth, POINT supports quick switching of HLS clients to their closest server, relying on its centralized topology management and the need to only update routing information at the entry points of the network. Since the number of servers needed depends on the streams served, production HLS services rely on multiple virtual servers which are activated and deactivated on demand. In an IP network, the HLS server is tied to a DNS name and an IP address, therefore switching to a different server either requires hiding all servers behind a load balancer (which is not easy if they are not co-located) or assigning multiple IP addresses to the same DNS name, and letting clients switch to an alternative IP address when their current one fails, a very slow process. In POINT, since names are used to locate content, a surrogate agent can instantaneously redirect requests from cNAPs to different sNAPs (and HLS server pools), transparently to the clients, by simply switching some path FIDs.




FIGURE 4 Logical trial topology.

3.3 | Service Deployment

Figure 4 shows the logical topology implemented for the trial. The shaded area in the figure represents the POINT network, which is connected via a set of NAPs to regular IP clients and servers. On the top of the figure, there are three servers, one offering the IPTV service over UDP/IP multicast, and two (primary and backup) offering the HLS service, connected to server side

NAPs (sNAPs). On the bottom there are a number of clients, each connected to a client side NAP (cNAP) via an ADSL network; the clients are either PrimeTel Set Top Boxes (STBs) for IPTV or laptops for HLS. The sNAPs and cNAPs are connected to two SDN switches, OVS1 and OVS2, interconnected via two trunk links (primary and backup). The IPTV server is also connected via two uplinks (primary and backup) to its sNAP. The SDN switches used *Open vSwitch* (OVS), controlled by an OpenDaylight controller (not shown), with a module implemented by the POINT team to bridge the ICN and SDN functionalities¹⁰. The POINT prototype used in the NAPs runs on regular Debian Linux 8. This topology is a simplified version of actual ISP topologies that span multiple cities: in each city there are one or more distribution networks, with a number of customers downstream and, possibly, some servers upstream. Normally, customers are served by local servers, but when such servers fail, they are served via the interconnection trunks by remote servers. It should be noted that the video servers are the production servers used by PrimeTel and the access network is PrimeTel's production ADSL network. The core network used for the trial is PrimeTel's R&D network. For comparison purposes, the POINT network runs side-by-side with a regular IP network with the exact same topology, using VLANs to allow both networks to operate all the way from the servers to the client devices.

Since the performance of PrimeTel's network and services is already of production quality, under normal operating conditions, POINT simply had to match this performance. The objective of the trial was rather to demonstrate that under exceptional network conditions POINT can result in a *better* experience for the viewer in terms of perceived QoE. For this reason, we subjected viewers to exceptional conditions with HLS and IPTV, both over the IP and the POINT network, and assessed both the objective performance of the network and the subjective evaluation of the service by the users.

For the HLS service, we used two bit rates with the same resolution but different qualities, as in PrimeTel's production service, albeit with a wider separation between the bit rates. We tested various bit rate combinations during a pilot study, and used a pair that were as close as possible but still led to a noticeable quality change. As an exceptional event, we brought the link to primary HLS server down for a few seconds. In the case of IP, the client timed out waiting for the primary server, before switching to the backup server. The HLS player in turn suffered from pixelation while waiting for the high-quality stream, before switching to the lower-quality stream. If the lower-quality stream timed out as well, playback froze. After bringing the link to the primary server back, the player eventually returned to the higher video bitrate. With POINT, clients switched automatically to the backup server, since it was the closest available to the client, without any noticeable effect to the end-users. When we brought back the link to the primary server, clients switched again to the closest available server, again with no effects on viewing experience.

For the IPTV service, we exploited the fact that there are two trunk links between the switches and two uplinks to the IPTV server. While in the IP case the spanning tree protocol uses only one link, with the other as a backup, in POINT both links can be exploited. In the IP case, we brought down the primary uplink to the server, which led to recalculation of the spanning tree and re-establishment of the IGMP snooping state, causing major viewing disruption. When the primary uplink was brought back up, the same steps were repeated, leading to another service disruption. With POINT, we brought down the primary trunk link between the SDN switches, which led to seamless switchover to the backup trunk link, while bringing back the primary trunk link led to another seamless switchover to that; in both cases, there were no noticeable effects on viewing experience.

4 | EXPERIMENTAL APPARATUS

4.1 | Goals and Hypotheses

Critical to the success of POINT is the *Quality of Experience* (QoE) of the video services, as perceived by the users¹⁹. QoE can be defined as a measure of users' quality needs and experiences arising from their interaction with an application, service, or system in a particular context²⁰. User QoE can be influenced by a number of factors, including technical, social and psychological, as shown in Figure 5^{18,21}. We therefore conducted a controlled (closed) trial of the POINT network to measure user QoE and usability of the video services offered. We decided to also measure usability, to determine whether the usability of the services had an impact on the perceived QoE of users. The trial took place in an operational ISP network (PrimeTel Cyprus) with users experiencing video content over both IPTV and HLS-based services. In the study, we considered the impact of the network (IP vs. POINT) and user context (individual vs. groups) on user QoE. We also measured usability in terms of *effectiveness*, *efficiency* and *satisfaction*²². The closed trial provided a plethora of insights about the performance of the POINT platform. These helped improve the POINT prototype and allowed its deployment in an open trial (at user homes) for an extended period of time.

The goal of our study was to evaluate user QoE with POINT in the presence of network congestion and link failures (deemed as exceptional operating conditions), and compare it to that of an ISP's conventional IP network. We evaluated user QoE and service usability using both subjective and objective metrics. We conducted the study in a set-up resembling an actual living

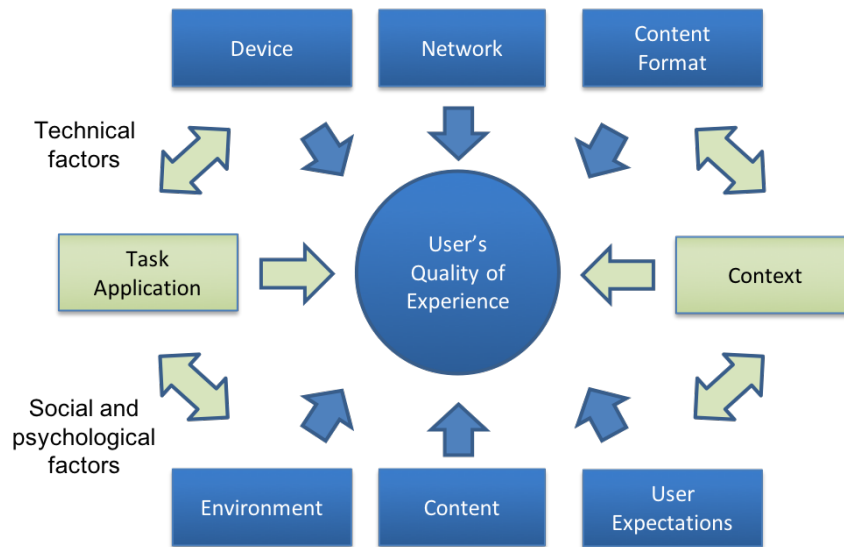


FIGURE 5 Factors that influence QoE (adapted from ¹⁸).



FIGURE 6 The area used for viewing with two participants.

room, shown in Figure 6. The area had space for three researchers to monitor and moderate each session. One of the researchers moderated the session and the second recorded various events of interest using the TechSmith Morae² software. The third researcher controlled the network and applied exceptional conditions on each type of network (IP and POINT). To minimize observer bias²³, only the lead researcher was aware of what was being studied. The rest of the team followed a study brief with step-by-step instructions on how to run the study. Also, none of the participants knew what was being studied, other than that the goal of the study was to study two prototype systems.

Participants experienced a variety of digital video content on a 48-inch TV terminal with a resolution of 1280 x 720 pixels (720p). The specifications of the video content used for each service are provided in Table 1.

TABLE 1 Characteristics of the video clips and Live TV channels.

Service	Compression Profile	Resolution	Video Bitrate	Framerate
HLS	MPEG-TS/H264 + AAC	720p	0.5+1.3 Mbps (two levels, variable)	25 fps
IPTV	MPEG-TS/MPEG-2/H264 + MP2/AAC	720p	3.5–5 Mbps (channel dependent)	25 fps

²<https://goo.gl/SgjP50>

To provide participants with a context of use for each service we created the following use case scenarios:

- *The HLS service offers movie trailers which are very popular in Cyprus. As we speak, there may be thousands of people watching these trailers at the same time as you. It is important to keep this in mind while you watch the trailers. In the first part of the study, you will watch two of the movie clips and in the second part of the study the remaining movie clips. Each clip lasts for 4 minutes on average.*
- *The IPTV service offers a higher quality of viewing experience than the HLS service. Imagine that your ISP has recently informed you about necessary network maintenance which could affect service delivery. It is important to keep this in mind when you watch TV. Please watch TV for 8 minutes in the first part of the study (any content that might be available) and another 8 minutes in the second part of the study.*

As a task, we defined viewing a video stream (a video clip with HLS or live TV with IPTV) from beginning to end over a network type (IP or POINT). Each task was repeated with both network types, but participants were not told which content was viewed over which network. During each task, we applied the same exceptional conditions to both network types. Specifically, for HLS we emulated high HTTP server latency by disconnecting the network link to the main HLS web server at regular intervals. For IPTV, we emulated link failure by disconnecting the primary link to the IPTV server or the primary trunk link at regular intervals. In both cases, POINT offered specific benefits. In the HLS scenario, POINT transparently exploited multiple servers offering the same content. In the IPTV scenario, POINT seamlessly switched to a secondary link when the primary link failed.

The following three hypotheses were tested in the study:


- H1:** The use of POINT results in better QoE for users of the HLS and IPTV services compared to IP. This is because POINT is more resilient under exceptional conditions, delivering video of better visual quality (e.g., without any major disturbances) compared to IP. These positive user experiences should be supported by relevant QoS data on the POINT network.
- H2:** Due to their minimalist design, the usability of the multimedia services (HLS and IPTV) is high. This means that participants are able to complete the assigned tasks (effectiveness) with little effort (efficiency) and high satisfaction. Therefore, their perceived QoE is not affected by the design of the services but only by the quality of the video streams.
- H3:** The presence of co-viewers increases the user QoE regardless of the network transmission mechanism used (IP or POINT). This is because watching videos with co-viewers increases each viewer's tolerance to artefacts or other visual elements that may be present in the multimedia stream as a result of applying the exceptional conditions.

4.2 | The Multimedia Services

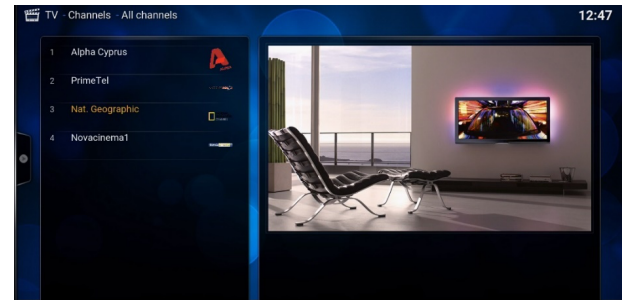
The HLS service offered four multimedia clips on a standard web interface, accessible using a standard web browser (see Figure 7a). Participants had to watch a different clip for each network type (IP or POINT). They could control video playback (e.g., pause, stop, volume, full screen) using a wireless mouse. The HLS webpage was served by NGINX and included the Video.js player with the adaptive bitrate HLS module and the movie trailer playlist. As a result, the player could decide on the quality of the video and adapt the playback based on the condition of the network. Each HLS video clip was encoded in 10 sec chunks with two video bitrates (see Table 1). This process created the main HLS playlist of each clip, which contained links to playlists for each different video bitrate. These playlists, in turn, described the bitrate, type and location of the video chunks.

The IPTV service consisted of a set of live TV channels delivered via the *Real-time Transport Protocol* (RTP) over UDP and IP multicast. The transmitter addressed each TV channel to a separate IP multicast address and the receivers subscribed to the address corresponding to the TV channel they wanted to watch. The service featured a simple menu from which users could select which channel to watch. Users were given a choice of four live TV channels (see Figure 7b). Users were asked to select any channel while experiencing content over both two network types (IP and POINT). They could also access other features of the IPTV service (e.g., settings) and control the playback (e.g., play and pause) of the stream using the TV's remote control.

To demonstrate surrogacy with HLS, we set up two identical (primary and backup) web servers on different machines. For the IP part of the trial, the IP clients resolved the DNS name of the HLS server to the IP of the primary server. Upon loading the webpage and clicking "play," the video player requested the higher bitrate playlist by default and began fetching the currently playing video chunk, as well as preloading the next 2-3 chunks. Whenever the exceptional conditions were activated (temporary loss of the primary HLS server), when the player tried to fetch the next chunks, the connection would timeout, switching the



(a) HLS



(b) IPTV

FIGURE 7 User interfaces of the video services.

player to the lower quality playlist; the player kept retrying the connection until it succeeded. When the primary HLS server became available again, the next lower quality chunk was fetched immediately, while for the next chunks, the player kept checking if it could upgrade to higher quality. For the POINT part of the trial, the DNS name resolved again to the IP of the primary HLS server, but POINT provided the added flexibility of introducing surrogate servers into the network. During the study, this functionality did not yet support HTTP in-session switching. Therefore the HTTP requests received by both web servers were monitored, and the backup server was introduced to the network only in between HTTP requests. This ensured that the users' viewing experience would not be interrupted when the same exceptional conditions as for IP were applied.

To demonstrate resilience to path failover with IPTV, two uplinks towards the IPTV server were created in the IP network. One uplink was designated as the primary, while the other one was designated as the backup. The primary uplink was normally in the forwarding state, while the backup uplink was normally in the blocking state. During the users' viewing session, the primary uplink was disconnected, triggering the recalculation of the spanning tree; eventually, the backup uplink was put in the forwarding state. Then the primary uplink was reconnected, which triggered the reverse procedure. In the POINT part of the trial, failover was demonstrated by establishing two trunk links between the SDN switches. The two trunk links were described in the POINT deployment configuration as a set of failover groups on each SDN switch. In this case, both trunk links were active, but only one was forwarding at any time. In the event of primary trunk link disconnection, the backup trunk link took over almost immediately. The reverse applied when the primary trunk link was re-connected.

4.3 | Participants

In total, thirty-nine users participated in this study. Participants were recruited randomly from various sources (e.g., Universities and the general public). Nine of the participants were used in a pilot study to evaluate the design of the experiment and the instruments of research. Those participants consumed the same video content as the others, but they spent overall more time in the laboratory to discuss improvements in the experimental set-up. The pilot study significantly contributed in finalizing the settings used in the main study. The remaining thirty participants were randomly assigned to the various conditions. Twenty-six of the participants consumed content in groups (Groups of viewers), and four individually (Single viewers). At the beginning of each session participants in groups were encouraged to interact with each other (e.g., express views about the content) while watching video. Participants were both male and female and with a variety of ages. Each group included a mixture of English (either native or as a second language) and native Greek speakers. The overall length of each testing session (individual or group) was approximately 40 minutes, including viewing time and answering questions.

4.4 | Experimental Conditions

The type of network transmission mechanism (IP vs. POINT) and service (HLS vs. IPTV) were investigated as within-subjects variables. The effect of social context (single viewers vs. groups of viewers) was investigated as a between-subjects variable. Participants were assigned at random in two groups (single viewers vs. groups of viewers) and to the eight experimental conditions shown in Table 2, which can be grouped in four sets (one per table row) where the same experiment was repeated over IP and POINT. In each experimental condition, we applied the same exceptional conditions to both the IP and POINT networks.

TABLE 2 Experimental Conditions.

Participants (N = 30)	IP	POINT
Single viewers	HLS Clips (1-4)	HLS Clips (1-4)
	IPTV Stream 1	IPTV Stream 2
Groups of viewers	HLS Clips (1-4)	HLS Clips (1-4)
	IPTV Stream 1	IPTV Stream 2

4.5 | Measures and Methods

We measured user QoE using an adapted version of the i-QoE questionnaire²⁴ and gathered a set of high-level QoS metrics which could potentially impact the user QoE (e.g., HTTP request rate). We measured usability as effectiveness, efficiency and satisfaction. *Effectiveness* was measured as the task completion rate. *Efficiency* was measured as the time needed to complete a task. *Satisfaction* was evaluated using the *Post-Study System Usability Questionnaire* (PSSUQ)²⁵. Finally, we observed and recorded several events of interest as users experienced the video content in each session. The events were defined prior to the experiment and were clustered in two categories: (1) User behaviour (e.g., feelings of frustration, viewer skips ahead, group interaction) and (2) System events (e.g., latency, blurriness). Table 3 shows these events.

TABLE 3 Events recorded in the study.

Event	Explanation
Interesting User Behavior	User behaviour of interest (e.g., viewer skips ahead, frustration)
Network Transmission	Network events of interest (e.g., perceived latency)
Video Artefacts	Video artefacts of interest (e.g., blurriness, blockiness)
Audio Artefacts	Audio artefacts of interest (e.g., background noise, synchronization))
Group Interaction	Group interactions of interest (e.g., chatting with each other)

We assessed the following *Objective Measures*:

- Completion rate: The completion rate was recorded as a binary measure of task success (coded as 1), task completed with difficulty (coded as 2) or task failure (coded as 0). The completion rate for each viewing experience was the number of users who completed the assigned task divided by the total number of users who attempted it. We captured the completion rate automatically using the TechSmith Morae platform.
- Task completion time: This is the average time (in seconds) of users completing tasks with each of the multimedia services. We captured the task completion time automatically using the TechSmith Morae platform.
- Session observation: An experimenter observed each user session and recorded events of interest (see Table 3) using the TechSmith Morae platform.
- QoS metrics: We gathered a very large amount of QoS data during the trial, but we only discuss below the QoS metrics we considered as important in the QoE assessment of the multimedia services (see Table 4).

We also assessed a set of *Subjective Measures* via two questionnaires:

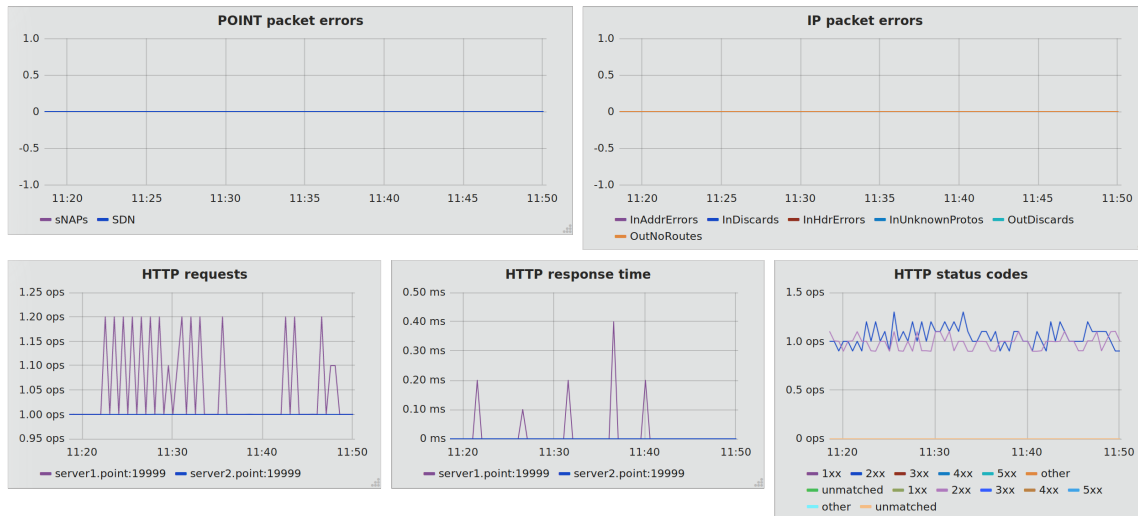
- The first questionnaire assessed the user QoE using the i-QoE questionnaire²⁴. The questionnaire includes items to measure the satisfaction, involvement, enjoyment, endurability and perceived visual quality of the digital video content. This questionnaire includes two items in a five-point Likert-scale (1=strongly disagree, 5 = strongly agree) and sixteen items in a seven-point Likert scale (1=strongly disagree, 7 = strongly agree). As the QoS of the network can vary from one multimedia clip to the other, so the user QoE can vary too. To capture these QoE variations, we asked participants to fill-in the i-QoE questionnaire after each viewing session (i.e. two clips or two TV channels), therefore it was administered twice.

TABLE 4 Some of the QoS metrics recorded in the study.

Name	Unit	Description
latency	ms	Average server-client latency
per-hop data rate	kbps	Average data rate through each network interface
HTTP request rate	sec ⁻¹	Primary and backup HTTP requests per second
HTTP response time	ms	Average primary and backup HTTP response time per request
time to repair a service	sec	(derived) Time required to recover from a failure

- The second questionnaire evaluated the usability of the services using the PSSUQ²⁵ standardised usability questionnaire. This questionnaire includes nineteen items in a seven-point Likert scale (1=strongly disagree, 7 = strongly agree). As this type of questionnaire has been designed to assess the overall satisfaction of users with aspects of the service (e.g., ease of use), it was administered only once, after users had completed all the assigned tasks with each service.

Finally, participants took part in a semi-structured interview. An experimenter led the session with a set of open-ended questions which gave participants the opportunity to provide more insights about the content they experienced with each service (e.g., if they noticed anything wrong with the video on each service). Also, in each session, the experimenter had a chance to ask participants some follow up questions to elicit additional responses. Each interview lasted 5-10 minutes.

**FIGURE 8** A network visualization panel.

4.6 | Monitoring the network

Since POINT is aimed at network operators, it must include monitoring and visualization tools to address operator needs. We have therefore created a monitoring library that reports statistics to a local monitoring agent in each POINT element. A monitoring server queries these agents to gather, log and visualize monitoring data. A rich set of monitoring points has been defined, which can easily be extended (and was, indeed, greatly extended for the needs of the trial). For example, the NAPs monitor IPTV related statistics such as received and transmitted bytes and channel acquisition times and HLS related statistics such as HTTP data rates and response times. In addition to POINT-specific tools, we used existing monitoring solutions to gather monitoring data from non-POINT devices, including STBs, servers and IP-based networking devices, such as operating system, per process/network interface and HTTP-level metrics. As an example, Figure 8 shows a visualization panel with data on network and application level performance over time, for example, packet errors at the POINT and IP levels, HTTP requests, response times and status codes. These tools were extensively used to gather the QoS data required to test our hypotheses during the POINT trial.

5 | RESULTS AND DISCUSSION

After the trial was complete, we processed the data gathered from TechSmith Morae, the questionnaires and the interviews, as well as the QoS metrics from the monitoring tools. This section presents the results of the QoE assessment part of the trial.

5.1 | Objective Assessment

5.1.1 | Completion Rate

Participants were able to complete all the assigned tasks successfully. This means that all participants watched content from beginning to end with both multimedia services and network types. However, participants using the POINT network had a more fluid viewing experience when the video was distributed over POINT, compared to when it was distributed over IP, with both the HLS and IPTV services. It is clear that the exceptional conditions had a bigger impact on the viewing experience of participants when the video streams were distributed on IP compared to POINT.

5.1.2 | Task Completion Time

Table 5 shows the time needed to complete the tasks. We can observe that the time to complete the IPTV tasks was nearly the same with both networks, but the time to complete the HLS tasks was higher with POINT. One possible explanation for this is that participants selected the slightly shorter video clips when using HLS over IP. However, the average time of usage for each service did not exceed 16 minutes on both network types (mean HLS = 14.8 and mean IPTV = 16). In conclusion, it is evident that participants completed the assigned tasks on both services (HLS and IPTV) with little effort.

TABLE 5 Task Completion Time.

Task	Total Time (sec)	Std. Deviation
HLS over IP	419.08	44.01
HLS over POINT	467.57	81.99
IPTV over IP	484.38	3.39
IPTV over POINT	485.93	7.76

5.1.3 | Distribution of Events

Table 6 shows the average number of interesting events recorded for each task (see Table 3 for the events recorded). A one-way ANOVA comparison taking the number of events as the dependent variable and the type of task as an independent showed that the number of events between the tasks significantly differed ($F(3, 40) = 4.306$; $p = .000$). A post-hoc comparison using Tukey's HSD test showed that the mean score for "HLS over IP" ($M = 4.55$, $SD = 1.51$) was significantly different than "HLS over POINT" ($M = 1.55$, $SD = 1.21$). It also showed that the mean score for "IPTV over IP" ($M = 6.91$, $SD = 2.02$) was significantly different than "IPTV over POINT" ($M = 1.36$, $SD = 1.29$). This shows that user sessions included many more events when participants watched video over IP compared to POINT. The most frequent events we recorded for video over IP were *Video artefacts* (e.g., *blurriness*, *blockiness*, *sharpness*, *clarity*) ($n = 27$) and *Network transmission* (e.g., *buffering*, *delay*, *frame loss*) ($n = 25$). An interesting observation was the way participants reacted when the artefacts occurred over IP. In the HLS service, the interruptions were short in duration, hence participants did not attempt to interfere with the video stream. However, it was easier for participants to become frustrated with the IPTV service. The interruptions generated stronger artefacts and had a longer duration. As a result, participants attempted several times to skip ahead or to change the channel. Overall, the distribution of events indicates that viewers had a more fluid viewing experience when watching video on POINT compared to IP.

5.2 | Subjective Assessment

Turning to the subjective metrics, we first analyzed the data from the pilot study. Our findings suggest that the network type did not have an impact on any of the questionnaire items. This shows that participants perceived the video streams on both network

TABLE 6 Distribution of Events.

Task	Number of Events	Std. Deviation
HLS over IP	4.55	1.51
HLS over POINT	1.55	1.21
IPTV over IP	6.91	2.02
IPTV over POINT	1.36	1.29

types as unimpaired. Then, in the main study we applied a series of one-way ANOVAs taking each questionnaire item as a dependent variable and the type of network transmission mechanism (IP or POINT) as the independent variable. We found a strong impact of the network type on multiple questionnaire items:

- Item 10 (“I was satisfied with the experience in watching this video clip”) ($F(1, 118) = 48.656; p < .001$)
- Item 11 (“I was pleased with the experience in watching this video clip”) ($F(1, 118) = 47.395; p < .001$)
- Item 12 (“I was contented with the experience in watching this video clip”) ($F(1, 118) = 35.349; p < .001$)
- Item 18 (“Did you perceive any visual impairments in the video (e.g., blockiness, blur, ringing”) ($F(1, 118) = 128.442; p < .001$)

TABLE 7 QoE Questionnaire Items Significantly Influenced by Network Type.

Item	POINT average	Std. Deviation	IP average	Std. Deviation
Item 10	5.60	1.09	3.82	1.65
Item 11	5.60	1.12	3.80	1.69
Item 12	5.38	1.24	3.78	1.68
Item 18	4.05	1.11	1.92	0.94

Table 7 shows the the detailed results for these items. The results for Item 10 suggest that users rated their viewing experience with both video services (HLS and IPTV) significantly more satisfying when the content was delivered over the POINT network compared to when it was delivered over IP. The pattern of increased satisfaction with the viewing experience over the POINT network is repeated in Item 11 and Item 12. It is evident that all participants found the experience of watching video to be more pleasing and felt more contented when the content was delivered over the POINT network compared to IP. The finding validates the hypothesis about POINT offering better QoE for users under exceptional conditions (see **H1** in Section 4), at least regarding the satisfaction of users. The average ratings for Item 18 show that participants thought that any visual impairments when video was distributed over the POINT network were imperceptible (and thus not annoying). On the other hand, they found the visual impairments when video was distributed over IP annoying. As a result, their viewing experience was more satisfying when the video was distributed on POINT as opposed to conventional IP.

We also examined the impact of social context (group vs. individual) on the user QoE. A Welch’s ANOVA test (unequal sample sizes) taking each questionnaire item as the dependent variable and the social context (group vs. individual) as the independent showed a significant effect of social context on item 15 (“The watching experience was rewarding”) (Welch’s $F(1, 19.50) = 6.96, p < .05$). Participants in groups perceived their watching experience as more rewarding (mean Group = 4.65) than individual participants (mean Individual = 3.56). A Two-Way ANOVA (adjusted for unequal sample sizes) did not show a significant interaction between type of network (IP vs. POINT) and social context (individual vs. group). These findings partially support the hypothesis about the presence of co-viewers increasing the users’ QoE (see **H3** in Section 4), as participants in groups had an overall more positive watching experience compared to individual participants. A possible explanation is that TV viewing created a social bond even when the participants were strangers^{26,27}. However, as this effect is independent of network type (IP vs. POINT), we can not attribute it to an increased tolerance to artefacts, as the presence of co-viewers did not change the participant’s ability to detect visual artefacts. These findings are further supported by the literature²⁴, which suggest the

presence of co-viewers has a positive impact on the enjoyment of the experience, but it does not have any impact on the ability of participants to detect visual artefacts. Overall, we are confident that the QoE differences are due to the exceptional conditions we applied on each network type and not due to a confounding variable (e.g., the presence of researchers), as this was the only aspect of the experimental design we significantly altered between the pilot and the main study.

Regarding the usability questionnaire, users rated the usability of both services overall high (HLS average = 5.07 and standard deviation = 1.19; IPTV average = 5.55 and standard deviation = 1.12). They also rated most qualities of the services high (see Table 8). The overall ratings for information quality may be related to the type of content participants experienced on both services: they may not have been interested enough in the content to rate it higher. A one-way ANOVA taking each questionnaire item as dependent variable and type of services as independent showed a significant effect for Item 2 ("It was simple to use this service") ($F(1,50) = 4.423$; $p < .05$). Participants thought it was significantly simpler to use the HLS (mean HLS = 6.42) than the IPTV service (mean IPTV = 5.96). This is to be expected, as the user interface of the HLS service is much simpler than that of the IPTV service. Therefore, we can conclude that participants' satisfaction with both services was high. Taken as a whole, the usability assessment findings (completion rate, task completion time and subjective satisfaction) support the hypothesis about the design of the services not having an impact on the viewing experience of participants (see **H2** in Section 4).

TABLE 8 Usability ratings for the two multimedia services (HLS and IPTV).

Usability Qualities	HLS	Std. Deviation	IPTV	Std. Deviation
System Usefulness	6.25	0.82	5.95	0.96
Information Quality	4.99	1.99	5.00	1.88
Interface Quality	5.49	1.38	5.49	1.38

5.3 | Post-Task Group Interviews

Finally, in the exit interview participants made several comments about their viewing experience. Overall, they thought that their viewing experience with POINT was better than with IP, regardless of the type of service they used. There were, however, some participants who did not notice any differences between the two networks when using the HLS service. As the exceptional condition we applied on HLS (HTTP server overload) resulted only in pixelation, rather than service interruption, it did not affect some participants as strongly as others. Below we present some representative comments about the participants' viewing experience with conventional IP and POINT.

- HLS over IP
 - "I wanted to stop the movie trailer and not watch it anymore."
 - "Yes, we could easily see the pixels. The resolution of the screen was not good. It was like 50 years ago."
 - "The quality of the two movie clips was not good. It was pixelized; the sound was not that good."
- IPTV over IP
 - "I noticed a lot of buffering, and it was unresponsive. It made us want to stop the TV."
 - "I noticed they were some interruptions."
 - "The picture froze very often."
 - "The channels stopped working and started buffering. Most likely there were too many people watching the same channel."
- HLS over POINT
 - "Watching the movie trailers, the second time (POINT) was more pleasing than the first (IP). The second set of clips was better than the first. However, the audio quality of both video clips was not better than the first set (IP)."
 - "Way better it was 60 fps; it was HD."

– “I did not notice any problems with the video feed. The quality was great.”

- IPTV over POINT

- “The second time it was better, a more pleasing experience. We did not have any buffering.”
- “The quality of the video was better. It was better than the first time.”
- “It was better than the first time; the picture did not freeze.”
- “It was OK the second time. I did not notice any problems.”

5.4 | QoS Assessment

To explain the subjective experiences of users, we looked into selected QoS metrics with both network types (POINT and IP). We present below the results of our analysis depending on the video delivery service.




FIGURE 9 HTTP data rate over IP.




FIGURE 10 HTTP response time over IP.

5.4.1 | HLS Service

To investigate the pixelization and buffering comments participants made for HLS over IP, a typical session was selected for analysis. The HTTP data rate for a typical session is shown in Figure 9, depicting traffic sent by the NGINX process on the primary and backup HLS servers. The maximum data rate was close to 18 Mbps, although this does not reflect the actual video bitrate, since HLS video is downloaded in pieces. We can observe that the exceptional conditions caused a disturbance in the data rate pattern. This is also reflected in very high HTTP response times, as shown in Figure 10, characteristic of an overloaded

server. Whereas normal response times were in the range of 100–400 ms, during exceptional conditions they exploded to up to 6 sec. This extra latency forced the client to automatically switch to a lower bitrate stream to lessen the load on the video server and avoid buffering and interruptions on the client end. Note that the client never switched to the backup server.




FIGURE 11 HTTP data rate over POINT.




FIGURE 12 HTTP response time over POINT.

Turning to HLS over POINT, the HTTP data rate of a typical user session can be seen in Figure 11. Although the peak data rate was again around 18 Mbps, the transmission was less bursty. With POINT, the introduction of the backup server into the network resulted in reassigning the clients to the closest server. Since this particular client was closest to the surrogate server, the loss of the primary server did not affect the user's viewing session as with HLS over IP. This can be seen in Figure 12, which shows that the client's requests are answered automatically by the backup server, with much lower latency than in the IP network (up to 250 ms vs. up to 6 sec). As a result, users did not experience any viewing disruptions when using HLS over POINT.

5.4.2 | IPTV

For the IPTV service, when we brought down the primary uplink to the server with the IP network, it led to recalculation of the spanning tree and re-establishment of the IGMP snooping state, causing major viewing disruption. When the primary uplink was brought back up, the same steps were repeated, leading to another service disruption. This can be seen in Figure 13, which shows the data transfer rates of the two uplinks from the IPTV server (bottom part) and of the downlink to the STB (top part). The user's viewing session begins with the first uplink being active (green line). Approximately at 15:01, the link is disconnected, which triggers the recalculation of the spanning tree and the re-establishment of multicast state, resulting in the user's video freezing (the black line drops to zero). Following convergence of the spanning tree, the second uplink becomes active (red line), and the video resumes at around 15:02:10. About 10 sec later, the first link is re-connected, which triggers another re-convergence event and video freezing (the black line drops to zero again). The first uplink becomes active at 15:03:20 and the

user's viewing resumes normally. Finally, another uplink disconnection at 15:05 resulted in another viewing disruption which was fixed at 15:06:10 through the aforementioned procedure.




FIGURE 13 Effect of link failover with IPTV over IP.

In the POINT case, the failure of the primary trunk link between the switches led to seamless switchover to the backup trunk link, while bringing back the primary trunk link led to another seamless switchover to that; in both cases, there were no noticeable disruptions in the service. This can be seen in Figure 14, which shows the data transmitted by the OVS switch on the left side of Figure 4 (ovs1, bottom part) and received by the OVS switch on the right side of the figure (ovs2, top part), while a user was viewing IPTV over POINT. Failure of the primary trunk link (eth2), which lasted around 20 sec, and fallback to the backup trunk link (eth1) are visible at 15:11 and 15:15. In contrast to Figure 13 and IP, recovery was almost instant, and the user's viewing session was not disrupted (the top line never drops to zero).




FIGURE 14 Effect of link failover with IPTV over POINT.

The total disruption time per failover event cannot be measured very accurately in the case of POINT, as the sampling rate of data acquisition (1 sec) is larger than the disruptions; based on the lack of any kind of effect on the user's screen, it was in the order of a few nanoseconds. On the other hand, the time needed to repair the IPTV service with IP ranged from 52 to 136 sec, with an average value of 59 sec, explaining the large gap between the user QoE assessment for POINT and IP with IPTV.

6 | CONCLUSIONS

Although implementing IP services over ICN is not a novel proposition, supporting any kind of IP-based service in an operator network differentiates POINT from other ICN projects, representing a novel way to introduce ICN to the Internet. The trial in PrimeTel's network was the first one for an ICN project, showing that the POINT prototype is stable enough to use on a real network, with actual services and unpredictable users.

As part of the closed trial, we conducted a set of carefully monitored experiments to evaluate the QoE offered by IP and POINT for video delivery services based on both HLS and IPTV. The analysis of the objective data found that the performance of the networks was similar under regular conditions. However, under exceptional conditions, such as link failure, POINT delivered a more fluid experience compared to IP. The finding is further supported by both the QoS assessment, which shows a marked deterioration of service delivery with IP, and the experiences of the viewers, as assessed by questionnaires and interviews. The subjective data shows a strong preference for video content distributed over POINT compared to IP, with participants finding the viewing experience with IP "annoying". This effect was particularly strong for IPTV as it produced a variety of different artefacts (e.g., frozen picture for several minutes, pixelization and sound distortion). However, the presence of co-viewers created a more positive watching experience for participants, regardless of the type of network used. This is to be expected, as people enjoy sharing a pleasant experience like watching a TV show. Finally, it is interesting that some participants thought that the quality of the video content they experienced was different on each network type. This further re-iterates the importance of a reliable networking delivery mechanism such as POINT. No matter how many pixels video specialists will be able to squeeze in a single packet, it is down to the networking delivery mechanism to ensure clients will experience it as intended.

Future work will investigate the impact of POINT on the user's QoE using *Electroencephalography* (EEG). This will allow us to gain deeper insights about the experiences of viewers while using each of these video services over the two network types.

ACKNOWLEDGMENTS

This research was supported by the EU funded H2020 ICT project POINT under contract 643990. We would also like to thank the members of the POINT team that participated in the design, implementation and testing of the POINT prototype used for the trials, including D. Trossen and S. Robitzsch (Interdigital), M.J. Reed and M. Al-Naday (University of Essex), G. Petropoulos, K.V. Katsaros and M.-E. Xezonaki (Intracom) and Y. Thomas and X. Vasilakos (AUEB).

References

1. Xylomenos G., Ververidis C.N., Siris V.A., et al. A Survey of Information-Centric Networking Research. *IEEE Communications Surveys Tutorials*. 2014;16(2):1024-1049.
2. Trossen D., Parisi G.. Designing and realizing an information-centric internet. *IEEE Communications Magazine*. 2012;50(7):60-67.
3. Jacobson V., Smetters D.K., Thornton J.D., Plass M.F., Briggs N.H., Braynard R.L.. Networking Named Content. In: Proc. of the 5th ACM Conference on Emerging Networking Experiments and Technologies (CoNext):1-12; 2009.
4. Trossen D., Reed M.J., Riihijärvi J., Georgiades M., Fotiou N., Xylomenos G.. IP over ICN - The better IP?. In: Proc. of the European Conference on Networks and Communications (EuCNC):413-417; 2015.
5. Cisco . *Visual Network Index: Forecast and methodology, 2016-2021*. Document ID:1465272001663118. Available at: <https://www.cisco.com/>; 2017.
6. Xylomenos G., Thomas Y., Vasilakos X., et al. IP over ICN goes Live. In: Proc. of the European Conference on Networks and Communications (EuCNC); 2018.
7. Xylomenos G., Phinikarides A., Doumanis I., et al. IPTV over ICN. In: Proc. of the ACM Packet Video Workshop; 2018.
8. Jokela P., Zahemszky A., Rothenberg C. Esteve, Arianfar S., Nikander P.. LIPSIN: line speed publish/subscribe inter-networking. In: Proc. of the ACM SIGCOMM Conference:195-206; 2009.

9. Reed M.J., Al-Naday M., Thomos N., Trossen D., Petropoulos G., Spirou S.. Stateless multicast switching in software defined networks. In: Proc. of the IEEE International Conference on Communications (ICC):1-7; 2016.
10. Petropoulos G., Katsaros K.V., Xezonaki M.E.. OpenFlow-compliant topology management for SDN-enabled Information Centric Networks. In: Proc. of the IEEE Symposium on Computers and Communications (ISCC):951-954; 2017.
11. Pantos R., May W.. *HTTP Live Streaming*. RFC 8216. Available at: <https://rfc-editor.org/rfc/rfc8216.txt>; 2017.
12. Al-Khalidi M., AL-Naday M., Doumanis I., et al. *Evaluation Report From System Trials*. POINT Deliverable D4.4. Available at: <https://www.point-h2020.eu/deliverables/>; 2017.
13. Fotiou N., Islam H., Lagutin D., Hakala T., Polyzos G.C.. CoAP over ICN. In: Proc. of the IFIP International Conference on New Technologies, Mobility and Security (NTMS):1-4; 2016.
14. Rahman A., Trossen D., Kutscher D., Ravindran R.. *Deployment Considerations for Information-Centric Networking (ICN)*. internet-draft: IRTF ICN RG; 2018. Available at: <https://datatracker.ietf.org/doc/draft-irtf-icnrg-deployment-guidelines/>.
15. Fenner B.. *Internet Group Management Protocol, Version 2*. RFC 2236. Available at: <https://rfc-editor.org/rfc/rfc2236.txt>; 1997.
16. Solensky F.T., Christensen M.J., Kimball K.. *Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches*. RFC 4541. Available at: <https://rfc-editor.org/rfc/rfc4541.txt>; 2006.
17. Dan A., Sitaram D., Shahabuddin P.. Scheduling Policies for an On-demand Video Server with Batching. In: Proc. of the ACM International Conference on Multimedia:15–23; 1994.
18. Timmerer C., Pereira F., Ebrahimi T.. *Quality of Experience in Multimedia Systems and Services: A journey Towards the Quality of Life*. ICME 2016 Tutorial; 2016.
19. Schatz R., Hoßfeld T., Janowski L., Egger S.. From Packets to People: Quality of Experience as a New Measurement Challenge. In: Biersack E., Callegari C., Matijasevic M., eds. *Data Traffic Monitoring and Analysis: From Measurement, Classification, and Anomaly Detection to Quality of Experience*, Berlin, Heidelberg: Springer 2013 (pp. 219–263).
20. Raake A., Egger S.. Quality and Quality of Experience. In: Möller S., Raake A., eds. *Quality of Experience: Advanced Concepts, Applications and Methods*, Springer 2014 (pp. 11–33).
21. Seufert M., Casas P.M., Wamser F., Wehner N., Schatz R., Tran-Gia P.. Application-layer monitoring of QoE parameters for mobile YouTube video streaming in the field. In: Proc. of the IEEE International Conference on Communications and Electronics (ICCE):411-416; 2016.
22. Nielsen J.. *Usability 101: Introduction to usability*. Available at: <http://www.nngroup.com/articles/usability-101-introduction-to-usability/>; 2003.
23. Sackett D.L.. Bias in analytic research. *Journal of Chronic Diseases*. 1979;32(1):51 - 63.
24. Zhu Y., Heynderickx I., Redi J.A.. Understanding the role of social context and user factors in video Quality of Experience. *Computers in Human Behavior*. 2015;49:412 - 426.
25. Lewis J.R.. Psychometric Evaluation of the Post-Study System Usability Questionnaire: The PSSUQ. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*. 1992;36(16):1259-1260.
26. Cohen L.E., Lancaster L.A.. Something to look forward to: Understanding the appeal of ritualistic television co-viewing events. In: Ryan K.M., Springer N.J., Noah J., Macey D.A., eds. *Friends, lovers, co-workers, and community: everything I know about relationships I learned from television*, Lexington Books 2016 (pp. 189–210).
27. Cooney G., Gilbert D.T., Wilson T.. The Unforeseen Costs of Extraordinary Experience. *Psychological science*. 2014;25(12):2259–2265.

