SDQ: Enabling Rapid QoE Experimentation using Software Defined Networking

Lyndon Fawcett[1], Mu Mu[2], Matthew Broadbent[1], Nicholas Hart[1], and Nicholas Race[1]

[1]School of Computing and Communications, Lancaster University, UK
[2]School of Science and Technology, The University of Northampton, UK

Abstract—The emerging network paradigm of Software Defined Networking (SDN) has been increasingly adopted to improve the Quality of Experiences (QoE) across multiple HTTP adaptive streaming (HAS) instances. However, there is currently a gap between research and reality in this field. QoE models, which offer user-level context to network management processes, are often tested in a simulation environment. Such environments do not consider the effects that network protocols, client programs, and other real world factors may have on the outcomes. Ultimately, this can lead to models not functioning as expected in real networks. On the other hand, setting up an experiment that reflects reality is a time consuming process requiring expert knowledge. This paper shares designs and guidelines of an SDN experimentation framework (SDQ), which offers rapid evaluation of QoE models using real network infrastructures.

Index Terms—Software Defined Networking; Quality of Experience; Experimentation; Adaptive Streaming

I. INTRODUCTION

High quality video streaming accounts for a large portion of today's Internet traffic [3], with HTTP adaptive streaming (HAS) established as the predominant method of streaming content across networks to heterogeneous devices. On its own, HAS offers improvements in Quality of Experience (QoE), as the video delivered is adaptive to the prevailing network conditions [6]. However, when other devices share the same network, fluctuations can occur [7]. Software Defined Networking (SDN) has opened up a wealth of opportunities for improving QoE [12], [6]. This is achieved through a greater level of control and awareness of the behaviour of the underlying network. However, one of the current issues within the QoE community relates to the popularity of simulation environments, particularly for the evaluation of QoE models [15] [13]. Simulations can overlook the effects of network protocols, client programs, or other real world factors and ultimately can lead to models not behaving as anticipated in real networks [5]. On the other hand, setting up an experiment environment that reflects real network configurations is a time consuming process requiring expert knowledge.

In this paper we go beyond the use of SDN to simply improve QoE, but consider how SDN can in fact support the rapid experimentation and testing of QoE models within a real network environment, thus addressing the aforementioned issue. We consider OpenFlow as a means to supporting innovation in QoE research in much the same way that it was originally conceived as a mechanism to enable researchers to innovate within networks. This paper highlights some of the challenges of integrating QoE models using hardware SDN equipment and shares experiences in realising a SDN experimentation testbed. We also demonstrate the benefit of SDQ using an existing QoE model [13].

II. BACKGROUND AND RELATED WORK

HTTP Adaptive Streaming (HAS) is widely adopted for video distribution. Using MPEG-DASH as an example, media content is encoded into *representations*, each of which is a version of the content prepared in *segments* using different encoding specifications. The adaptation logic (between representations) is often determined at the player, with decisions based upon criteria such as estimated throughput [9] and buffer occupancy [8]. There has been significant effort placed in improving the QoE of adaptive video streaming. Most of this work in this area including [7], [8], [10] focuses on optimising the network efficiency or the QoE on individual media streams. Without coordination between HAS clients, TCP-based resource allocation can lead to unfairness at the user level [13] and severe fluctuations in multi-stream environments [10]. With the increasing number of HAS streams, it is essential to orchestrate network resource consumption through 1) a better understanding of the user-level requirements of user applications, and 2) a transparent network resource allocation service in content networks.

Software Defined Networking [16], through the decoupling of the control plane from the packet forwarding plane, allows network services such as load balancing and context aware resource allocation to be realised and automated as part of the service delivery chain [12]. This has led to a significant body of research investigating the use of SDN for QoE assessment and improvement including [11], [4]. We first demonstrated the feasibility of a SDN-assisted video quality management framework in [6]. Through fairness modelling using video quality, switching impact and cost efficiency as the impact metrics [14], a scalable resource allocation model *UFair* was introduced. This is designed to improve delivered video quality and user-level fairness between HAS media streams [13].

QoE models are often proven using programmatic simulation for their ease and rapid nature. Simulations overlook some networking and client aspects, which result in models not working as expected in real networking environments [5]. In contrast, to test a QoE model in a real-world environment requires multiple switches and even more clients, all of which

consume more resources and time than a simulation. The network emulator Mininet is often considered as an appropriate solution to the aforementioned problems; it is capable of running at large scale, uses a real networking stack, and can execute client programs [1]. However, at the moment it is limited by the underlying capabilities that the software switch (Open vSwitch) provides. In particular, Open vSwitch is missing some of the extended features of OpenFlow 1.3 that are important for QoE, such as *metering* (rate-limiting of flows). This highlights the need for a system that can balance these requirements: an SDN and virtualisation environment that provides the realism of real world experimentation whilst offering the benefits in agility and low cost associated with simulation environments.

III. SDQ Framework

In the following section, we describe the requirements and architecture of the *SDQ Framework*; a harness to be used in aiding rapid deployment and orchestration of experiments. As such, the architecture and experimentation environment has to fulfil the following requirements:

- *Experiments Close to Practice and at Scale*: To provide both realism and scale, the environment should encompass both physical and virtual elements.
- *Software Defined Dynamic Manipulation of Available Bandwidth Within the Network*: To match real-world network topologies, the link speeds must be rate limited. For QoE enforcement, dynamic configuration of rate limiting also needs to be applied to specific flows.
- *Configurable Clients*: The client's configuration should be changeable between experiments.
- *Rapid Repeatability of Experiments in a Clean Environment*: Ensure that no residual effects are left over from previous experiments.
- *Generic Framework*: The system should offer control and statistics from physical networks and any virtual infrastructure.

The SDQ framework orchestrates the virtual and physical network infrastructure using SDN to assist the execution and statistical data gathering of network based QoE experiments. It consists of a three layer architecture: the top layer contains components provided by the researcher including the test manifest and application/user-level functions such as a QoE module. The middle layer contains the SDQ orchestrator which interfaces with, and includes, the infrastructure managers. The bottom layer contains the network and virtualisation infrastructure where the experiments are deployed. The following is a breakdown of the components shown in Figure 1:

Test Manifest: describes the experiment in a JSON format. It includes each of the clients' IP addresses, the networks each is attached to, the virtual machine image to be used, network emulation requirements, and timestamps for automated tests.

QoE Model: an interchangeable component which communicates with the SDQ orchestrator (described below) through an RPC (Remote Procedure Call) interface providing information about resource allocation of one or more media flows.
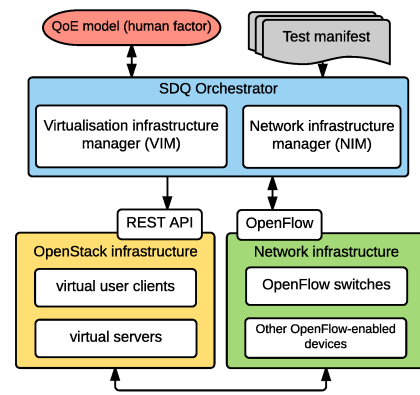

Fig. 1. SDQ framework

Additionally, information is sent back in regards to the current throughput at different points in the network using OpenFlow's *meter statistics* and *flow statistics* messages.

SDQ Orchestrator: handles communication between all of the components. It includes two subcomponents, the *Virtual Infrastructure Manager (VIM)* and *network Infrastructure Manager (NIM)*. The VIM controls the virtualisation infrastructure through a RESTful API, it launches and configures experiment nodes with information from the test manifest. At the end of the experiment it resets the test environment, removing networks and virtual machines it instantiated, so that the environment is ready for the next experiment. The NIM controls the network infrastructure and consists of an OpenFlow controller containing a metering and monitoring application. It installs meter flow mods on request from the QoE model and provides information from the network including current throughput of flows and switches.

The virtualisation infrastructure (through integration with OpenStack) instantiates experiment nodes and exposes them to the network. The network infrastructure creates connections between nodes and switches and provides a platform for QoE-based flow enforcement. Section IV provides further details regarding the configuration of each of these elements.

IV. Experimentation environment

This section describes the virtualisation and network infrastructure used to create the environment for SDQ.

Our objective is to create network and virtualisation infrastructures that are controllable through the SDQ framework[1]. Ultimately, these need to resemble real world deployments. To illustrate the advantages of SDQ, we highlight a potential deployment scenario for evaluating QoE. This consists of a number of households connected to the Internet using consumer ADSL connections. They all share a common DSLAM, which is connected to the wider Internet over a restricted shared link. The video servers are located at the remote end of this connection. The home routers and the local DSLAM are also under SDN control to provide link emulation and QoE enforcement functionalities. To realise the topology described above we create the experimentation environment shown in
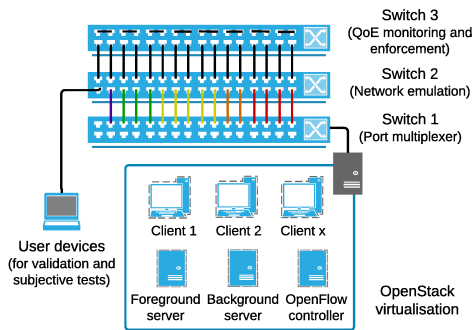
[1]https://github.com/LancasterNetworking/SDQ

Fig. 2. Experiment environment

Figure 2. It consists of a large cluster of generic servers, three OpenFlow-capable switches, a user device, and local connectivity between each.

In the following subsections, we outline the process undertaken to establish our experimental environment so that readers can replicate our experiments.

A. Virtualisation Infrastructure

At the core of the virtualisation infrastructure is an OpenStack installation. This provides the means of building and connecting virtual machines (VM). The OpenStack installation is standard, with one main modification: VLAN trunks are used to break-out network interfaces from virtual machines. These are then mapped one-to-one to exclusive physical interfaces on a switch. This is an essential feature for our experimentation as it allows each client to be directly assigned to a physical port on an SDN controlled switch.

A typical experiment requires 10s to 100s of OpenStack configuration elements, most of which are replicas of a basic template. Building topologies by hand, either via the *Horizon* GUI, or using command line, is tedious and prone to error; the ability to both build and destroy test topologies rapidly and consistently is essential both for replicating, revisiting or extending experimental results and for sharing the test environment with other workers. As such, we developed an orchestration tool titled *MiniStack*[2], which takes topology description files and builds and boots fully connected experimental topologies. The tool is written in Python and uses the OpenStack REST API.

B. Network Infrastructure

The network infrastructure used in this example consists of two OpenFlow v1.3 capable switches (Switch 2 and 3 shown in Figure 2) with metering support. In our facility, we use Hewlett Packard Enterprise's HP3800 switches, as they fulfil both of these requirements. However, other compliant switches could be used instead. The HP3800 also hosts other important capabilities, such as the ability to flexibly partition a single physical switch into a number of virtual switches. Each of these is a complete, distinct OpenFlow instance. This too is outside of the scope of OpenFlow, but is a feature present on a number of devices available on the market. This

[2]https://github.com/hdb3/ministack

partitioning feature is vital in achieving the scale required in experimentation without incurring the associated cost.

The network infrastructure is controlled by the NIM: a Ryu [2] controller application *OpenFlow Bandwidth* that provides a REST/JSON API to issue requests for bandwidth management and monitoring using simple intuitive JSON defined messages. Ryu was selected because of its support for extended OpenFlow 1.3 features and support for various hardware switches. For the application, a typical request would be to report the current network traffic level for a port or previously defined flow.

Overall, the combination of features, programmability, and openness provided by OpenFlow greatly assist us to fully realise QoE applications in real-world networks.

V. QoE MODEL EXPERIMENTATION

This section describes how the SDQ experimental environment (described in Section IV) was used to evaluate the UFair QoE model [13], which has previously been evaluated via means of a simulator. When transforming this QoE model into a real networked implementation, we came across a number of issues. During early testing HAS clients were not receiving sufficient bandwidth to reach the stream quality that the model was aiming to achieve. Following analysis, the cause was determined to be an assumption made in the simulator that if meters were set to a specific bandwidth then clients would receive exactly that. In addition, the simulator did not consider packet header sizes, the meter dropping policy, and that HAS chunks can vary in size. These were factors that were easy to overlook when creating a simulation. In the remainder of this section, we describe in more detail the experimental setup, the experiments performed and the results of doing so.

A. Experiment topology

The topology used for the evaluation of UFair was consistent with the original simulated environment. It represents a tiered multi-household network, in which each household contains 4-6 hosts running DASH clients, all of which are connected to a gateway. This gateway is then connected, along with other gateways in the topology, to an aggregation switch. Over another hop, a foreground (serving DASH content) and a background server act as endpoints for their respective traffic.

To simulate a potential home network environment where links are limited, emulation of network link characteristics are used. Through SDQ, the links between access switches and the gateway switches are limited to 20Mbps. Similarly the link between the aggregation switches is restricted to 50Mbps. The sum of the connectivity available to household links is 100Mbps, This results in a situation whereby there is more demand than there is supply in the case of multiple households. In these circumstances, the adaptive streams in each house are affected by hosts within the same house, as well as the behaviour of hosts in other houses.

B. Experiments

In order to evaluate the effectiveness of the UFair QoE model, and to assess the capability of SDQ to achieve the
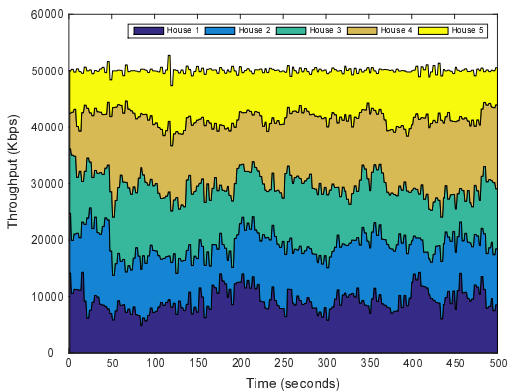
Fig. 3. No OoE model across households

Fig. 4. UFair across households

required functionality, we conducted two representative experiments: one where the QoE model was used for network resource allocation and the other with no QoE model applied. SDQ was used to capture experiment data, monitor traffic, and enforce traffic limits per flow and per household.

C. Results

This section compares the results of the two experiments (one with QoE model active and the other without) ran with the assistance of the SDQ framework. Figures 3 and 4 depict the allocation of network resources on the aggregation link between all households. Without a global network view, TCP-based resource allocation causes fluctuations in the network (Figure 3), which ultimately deteriorates the user experience through an increased switching of streams.

Through this experimentation, SDQ has shown to be an effective framework to support QoE experiments relating to adaptive video content. The whole toolset, including streamlining the orchestration of the QoE model, virtualisation infrastructure, and physical OpenFlow equipment, allows researchers to focus on human factor modelling and helps to verify the feasibility of any QoE model.

VI. CONCLUSIONS

Often, researchers use simulations to test QoE models due to the ease, agility and low cost that they offer when compared to real world testing. However, simulations often fail to recognise some of the additional effects that are present in actual networks and the technologies that use them. This leads to models behaving differently in reality, lessening the contribution of experimental findings. This paper introduces SDQ, a framework that uses SDN to facilitate rapid experimentation of QoE models in such a realistic environment. SDQ aids QoE researchers by reducing the barrier to entry for SDN-assisted QoE experiments. An example use case, using the UFair model, demonstrates how a QoE model previously tested solely in simulation can be evaluated in the real world.

REFERENCES

[1] Mininet network emulator. *http://mininet.org/*.
[2] Ryu OpenFlow controller. *http://osrg.github.io/ryu/*.
[3] Cisco Visual Networking Index: Forecast and Methodology, 2014-2019 White Paper, 2015.
[4] P. Agyapong, M. Iwamura, D. Staehle, W. Kiess, and A. Benjebbour. Design considerations for a 5G network architecture. *Communications Magazine, IEEE*, 52(11):65–75, 2014.
[5] S. Floyd and V. Paxson. Difficulties in simulating the internet. *IEEE/ACM Transactions on Networking (TON)*, 9(4):392–403, 2001.
[6] P. Georgopoulos, Y. Elkhatib, M. Broadbent, M. Mu, and N. Race. Towards Network-wide QoE Fairness Using Openflow-assisted Adaptive Video Streaming. In *ACM SIGCOMM 2013 Workshop on Future Human-centric Multimedia Networking (FhMN)*, 2013.
[7] T.-Y. Huang, N. Handigol, B. Heller, N. McKeown, and R. Johari. Confused, Rimid, and Unstable: Picking a Video Streaming Rate is Hard. In *Proc. ACM IMC*, pages 225–238, 2012.
[8] T.-Y. Huang, R. Johari, N. McKeown, M. Trunnell, and M. Watson. A buffer-based approach to rate adaptation: Evidence from a large video streaming service. In *Proceedings of the 2014 ACM conference on SIGCOMM*, pages 187–198. ACM, 2014.
[9] J. Jiang, V. Sekar, and H. Zhang. Improving Fairness, Efficiency, and Stability in HTTP-based Adaptive Video Streaming with FESTIVE. In *Proc. ACM CoNEXT*, pages 97–108, 2012.
[10] Z. Li, X. Zhu, J. Gahm, R. Pan, H. Hu, A. Begen, and D. Oran. Probe and Adapt: Rate Adaptation for HTTP Video Streaming At Scale. *IEEE Journal on Selected Areas in Communications*, 32(4):719–733, 2014.
[11] E. Liotou, G. Tseliou, K. Samdanis, D. Tsolkas, F. Adelantado, and C. Verikoukis. An SDN QoE-Service for dynamically enhancing the performance of OTT applications. In *Quality of Multimedia Experience (QoMEX), 2015 Seventh International Workshop on*. IEEE, 2015.
[12] R. Mijumbi, J. Serrat, J. L. Gorricho, N. Bouten, F. De Turck, and R. Boutaba. Network Function Virtualization: State-of-the-art and Research Challenges. (c):1–28, 2015.
[13] M. Mu, M. Broadbent, N. Hart, A. Farshad, N. Race, D. Hutchison, and Q. Ni. A Scalable User Fairness Model for Adaptive Video Streaming over Future Networks. *IEEE Journal on Selected Areas in Communications*, 2016.
[14] M. Mu, S. Simpson, A. Farshad, Q. Ni, and N. Race. User-level Fairness Delivered: Network Resource Allocation for Adaptive Video Streaming. *2015 IEEE/ACM International Symposium on Quality of Service (IWQoS)*, 2015.
[15] M. Seufert, T. Hosfeld, and C. Sieber. Impact of intermediate layer on quality of experience of HTTP adaptive streaming. *2015 11th International Conference on Network and Service Management (CNSM)*, 17(1):256–260, 2015.
[16] W. Xia, Y. Wen, C. H. Foh, D. Niyato, and H. Xie. A Survey on Software-Defined Networking. *IEEE Communications Surveys & Tutorials*, 17(1):27–51, 2015.