Demonstrating the Prospects of Dynamic Application-Aware Networking in a Home Environment

Florian Wamser, Thomas Zinner, Lukas Iffländer, Phuoc Tran-Gia
University of Würzburg, Institute of Computer Science
{wamser,zinner,ifflaender,trangia}@informatik.uni-wuerzburg.de

Categories and Subject Descriptors

C.2.1 [**COMPUTER-COMMUNICATION NETWORKS**]: Network Architecture and Design—*Network communications*

Keywords

Application-Aware Networking, Dynamic Resource Allocation, Home Networks, YouTube

1. INTRODUCTION

The success of tablet computers, game consoles, and Smart TVs reflects the increased user demand for Internet-based services at home. The users in the home network can access value-added services offered directly by the network provider, such as IPTV. Likewise, they also use Over-The-Top (OTT) services like YouTube, Netflix, or online gaming and browse the web or download files. All these services have specific requirements with respect to the network resources which have to be fulfilled to ensure a good Quality-of-Experience (QoE) for the users. Furthermore, multiple users may concurrently access different services via the central Internet access point in the home network, the home gateway, cf. Figure 1.

As stated by the Home Gateway Initiative (HGI) [1], the network at home and the broadband Internet access link may constitute a bottleneck. This may be due to the insufficient availability of broadband access, i.e., the network provider offers only smallband Internet access, or due to limitations within the home network, like the varying channel quality of the wireless networks.

Traffic in today's network structures is usually transmitted on a best effort basis. As a result, different services or applications with varying requirements and capabilities are treated equally on a per flow-basis, resulting in a unfairness in terms of QoE. Long lasting OS updates on a home computer may thus interfere with video streaming to a Smart TV which leads to video stallings and a degradation in the QoE of the user. In such a case, it is necessary to explicitly allocate the network resources unequally to the involved applications on network level to achieve a similar QoE across multiple applications.

SIGCOMM'14, August 17–22, 2014, Chicago, IL, USA.
ACM 978-1-4503-2836-4/14/08.
http://dx.doi.org/10.1145/2619239.2631450.

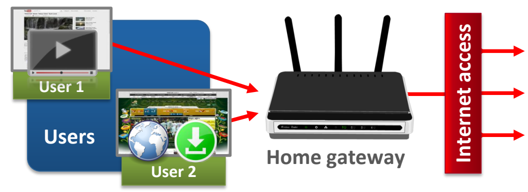

Figure 1: Illustration of a home network.

This explicit network resource control is known as Application-Aware Networking. Scalability issues hinder the implementation of such mechanisms within aggregation and wide area networks, but the small size of the home network makes it a promising candidate. Most of the traffic is forwarded in home networks via the home gateway, making it possible to control the network resources, and therewith the application quality, at this entity.

In this demonstration, we show the potential of flexibly allocating network resources to different applications. We focus on a two-application scenario where YouTube flows and a file download compete for resources via a shared bottleneck link. YouTube maintains a playback buffer to overcome resource limitations on a short time scale. We take advantage of this in order to provide an accurate reaction against video stallings. We implemented a network-based buffer estimator for YouTube which allows the accurate monitoring of the application buffer state. If the local video buffer runs empty, the IP flow is prioritized, if the buffer is sufficiently filled, the prioritization is turned off, or other applications like browsing are prioritized.

2. IMPLEMENTATION

This section highlights details of the monitoring and resource allocation mechanisms implemented on the home gateway, an AVM FRITZ!Box 3390 [2].

The user perceived quality of a service depends on a multitude of technical and non-technical factors. Some of the technical factors are application-specific indicators for the QoE, like the used video codec, the filling of the playback buffer, or the perceived delay on application level. For a targeted resource allocation in the network an accurate monitoring of the application state, i.e. of these indicators, is necessary. For YouTube video streaming, the buffered play-

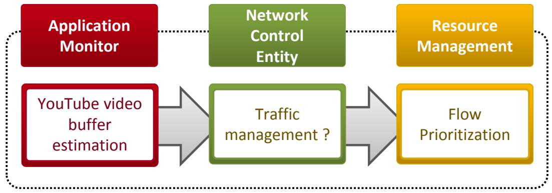

Figure 2: Control loop taking application information into account.

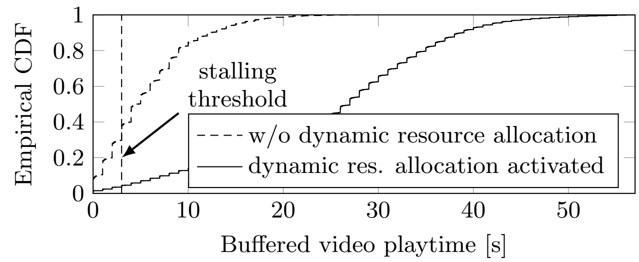

Figure 3: CDF for YouTube buffer level.

time is such a state parameter which has a high impact on the QoE [3]. If the buffer falls below a certain threshold, the video stalls and the QoE is affected.

To demonstrate a dynamic resource allocation of network capacity for two applications, namely YouTube video streaming and a parallel download, we implemented the necessary functionality on a FRITZ!Box running Freetz OS [4]:

(a) *YouTube video buffer estimation*, (b) *flow prioritization / flow de-prioritization* as resource management actions, and (c) *network control entity* that decides what value of application status leads to which resource management action, see Figure 2.

The resource management process is as follows: The application monitor detects a YouTube flow based on the HTTP GET URL of the YouTube video request [5]. For HTTP Secure (HTTPS) requests, a trusted man-in-the-middle proxy is installed. If a YouTube video and its corresponding IP flows are detected, the application monitor estimates the buffered playtime of the client by inspecting relevant data parts in the IP packets . Afterwards, the IP flow signature and the estimated buffered playtime is signaled to the network control entity. It decides whether the IP flows are prioritized or de-prioritized. The prioritization is performed with traffic control (tc) by setting up classful queueing disciplines in the standard Linux kernel of the home gateway.

To illustrate the prospects of this dynamic prioritization, we present results for an exemplary scenario. It consists of one YouTube video competing for bandwidth in a 3 Mbit/s IEEE 802.11b/g wireless access network with two file downloads. First, the YouTube video is started. Second, the downloads are started after 15 s. The YouTube video has an average video encoding rate of 2.9 Mbit/s.

Figure 3 shows the cumulative distribution function (CDF) of the buffered playtime in seconds at the client. The dashed line on the left in the figure corresponds to the stalling threshold where the playback at the YouTube player is restarted after buffering due to buffer underrun. In particular, this means that the buffer level should always exceed this threshold to avoid stalling and ensure a good QoE for this video [3]. In the first case, the buffered playtime is depicted without the proposed dynamic resource allocation. In 40 % of the time, the buffer is under the stalling threshold resulting in a bad QoE. In contrast, with the resource allocation activated, the buffer level can be maintained at a high level.

3. DEMONSTRATION SCENARIO AND PRESENTATION

To demonstrate the functionality of the dynamic resource allocation at the home gateway, as well as to foster an active discussion with the audience, we stick to the same scenario as described in the last section. Two Laptops are connected to the home gateway via wireless LAN, serving as end users which run the applications. Additionally, a detailed graphical user interface for YouTube has been developed to illustrate the functionality and to allow for a good presentation. As a consequence, viewers of the demonstration can follow the course of the buffer level of YouTube and see when the playback behavior is affected and when not.

We plan to show a predefined demonstration sequence, but any YouTube video or download can also be started manually at the request of viewers. The predefined sequence is as follows: (1) User 1 starts to watch a YouTube video. (2) After 15 s, User 2 is starting some downloads which overwhelm the access network. Consequently, the YouTube video experiences problems and the buffer level is decreasing, until the video stalls for several times. (3) Afterwards, we activate the resource allocation which detects the YouTube buffer level and adjusts the prioritization for the flows. In the end, this allows YouTube to download video data at the expense of the bandwidth for the downloads. The audience can see that YouTube is able to load the buffer up to a certain buffer level, while the downloads are automatically throttled if YouTube is below the particular buffer level. Consequently, the perceived quality of the YouTube user is good since stalling does not occur anymore. The downloads are throttled. From a QoE perspective, this however only affects the downloads to a minor extent.

That way, the demo's audience can observe the influence of the resource management mechanism for YouTube on the situation in the network as well as the improvements in QoE for the users. A screencast of this demonstration is available at `https://www.youtube.com/watch?v=Nk2Pjht6YVk`.

4. REFERENCES

[1] Home Gateway Initiative, "Home Gateway QoS Module requirements," whitepaper, Dec. 2012.

[2] AVM, "FRITZ!Box 3390," http://www.avm.de/en/Produkte/FRITZBox/FRITZ_Box_3390/.

[3] T. Hoßfeld, R. Schatz, E. Biersack, and L. Plissonneau, "Internet Video Delivery in YouTube: From Traffic Measurements to Quality of Experience." Springer's Computer Communications and Networks, 2013.

[4] Freetz Project, "Firmware image for home routers," http://freetz.org/.

[5] F. Wamser, D. Hock, M. Seufert, B. Staehle, R. Pries, and P. Tran-Gia, "Using Buffered Playtime for QoE-Oriented Resource Management of YouTube Video Streaming," *Transactions on Emerging Telecommunications Technologies*, vol. 24, Apr. 2013.