

Implementing Application-Aware Resource Allocation on a Home Gateway for the Example of YouTube

Florian Wamser^(✉), Lukas Iffländer, Thomas Zinner,
and Phuoc Tran-Gia

University of Würzburg, Am Hubland, Würzburg, Germany
{wamser,ifflaender,zinner,trangia}@informatik.uni-wuerzburg.de

Abstract. Today’s Internet does not offer any quality level beyond best effort for the majority of applications used by private customers. If multiple customers with heterogeneous applications share a bottleneck link to the Internet, this often leads to quality deterioration for the customers. This particularly holds for home networks with small band Internet access and for home networks with resource limitation like a bad channel quality within a wireless network. For such cases, the best effort allocation of resources between heterogeneous applications leads to an unfair distribution of the application quality among the users. To provide a similar application quality for all users, we propose to implement an application-oriented resource management on a home gateway. Therefore, allocation mechanisms need to be implemented such as the prioritization of network flows. Furthermore, a component monitoring the application quality and dynamically triggering these mechanisms is required. We show the feasibility of this concept by the implementation of an application monitor for YouTube on a standard home gateway. The gateway estimates the YouTube video buffers and prioritizes the video clip before the playback buffer depletes.

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.

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 small band 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, as introduced in [2].

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 them 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. The applicability of this approach was shown in [3], however, performance studies on the potential of the approach on a home gateway are missing.

In this paper, we evaluate 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 video pre-buffering 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.

The remainder of this paper is structured as follows. Section 2 summarizes the home gateway architecture and its components. In Sect. 3, the implementation details for the specific scenario are described. Section 4 highlights the evaluation setup, and Sect. 5 presents the results of our evaluation. The related work is summarized in Sect. 6, and the paper is concluded in Sect. 7.

2 Application-Aware Resource Control for a Home Gateway

In this section, we introduce the developed application-aware resource management framework running on the home gateway and its components.

The developed architecture consists of four components. These are a network monitoring, a network control component, an application monitor, and a decision component. The purpose of the architecture is a cross layer approach that manages the network resources in a more sophisticated way in order to increase the QoE of all network users. It manages the network resources based on the application information. Resource management actions are only triggered if they may improve the average application performance of all network users.

This means that they are triggered if an indication for, or the QoE degradation itself, has been detected and a resource management action may improve the QoE of one application while not degrading the quality of others. In the following, we describe the four components.

Network Monitoring. The network monitoring and control component firstly performs network monitoring of relevant QoS parameters and keeps track of the current network load. It provides a collection of relevant QoS parameters to the decision component. Thus, it supports the decision component by forcing appropriate resource management actions.

Network Control. This component implements possible resource control mechanisms for an explicit allocation of network resources to single flows or traffic classes. With respect to the capabilities of the home gateway, these mechanisms may range from additional entries in the routing table to the prioritization of application classes or flows, sophisticated resource allocation actions or even an active manipulation of the application layer content. The network control component implements the input from the decision component.

Application Monitor. The application monitor continuously estimates the current application quality based on the applications' network flows. The quality depends on the current application state which is a collection of application-specific quality indicators. These quality indicators describe whether the current resources offered by the network are sufficient for a good quality. The resulting application quality is forwarded to the decision component.

Decision Component. The decision component sequentially determines how the available resource control mechanisms shall be forced based on the current network and application monitoring data. It predicts how the resource control may change the application and network state. A resource control action is triggered if the confidence in the prediction is high, the degree of the QoE improvement is significant, and the effort performing the resource management action is acceptable.

3 Implementation Details and Application Design

This section highlights the implementation details of our approach.

3.1 Monitoring the YouTube Buffer State

Details of our approach to estimate the application quality of YouTube are discussed in this subsection.

To enable a smooth playback of YouTube videos for the customer with a fluctuating network bandwidth, the YouTube application buffers a certain amount of the video data. To further achieve an economically efficient service, the pre-buffered video data have to be kept as small as possible to minimize the waste of

transport resources if the user aborts the video playback. This task is performed by different flow control algorithms. Thus far YouTube implements two different approaches on flow control: the “throttling approach” performed on the server side by controlling network bandwidth [4] and the “range request approach” performed by the client by requesting the required parts of the video similar to the MPEG-DASH approach [5].

YouTube supports multiple media container formats [6]. Each format encapsulates the content provided in a different way. The most common formats are Flash Video (FLV) [7], Third Generation Partnership Project (3GPP) [8] and MPEG 4 (MP4) [9] where MP4 can be implemented as a continuous or fragmented format, the last one currently being the default format for newly uploaded videos to the platform.

MP4 specifies that media and meta data can be separated and stored in one or several files, as long as the meta data is stored as a whole. The file is divided in so-called boxes. The standard files contain one “ftyp”, “mdat” and “moov” box. Where “ftyp” specifies the file type, “mdat” contains the media data and “moov” the meta data. “moov” has several subboxes that contain the video header and timescale, information about substreams and the base media decode time.

Fragmented MP4 provides the ability to transfer the “mdat” box in multiple fragments instead of one big blob. As shown in Fig. 1 each fragment consists of a “moof” and a “mdat” box. The “moof” (movie fragment) box contains subboxes determining the default sample duration, the number of samples in the current track and a set of independent samples.

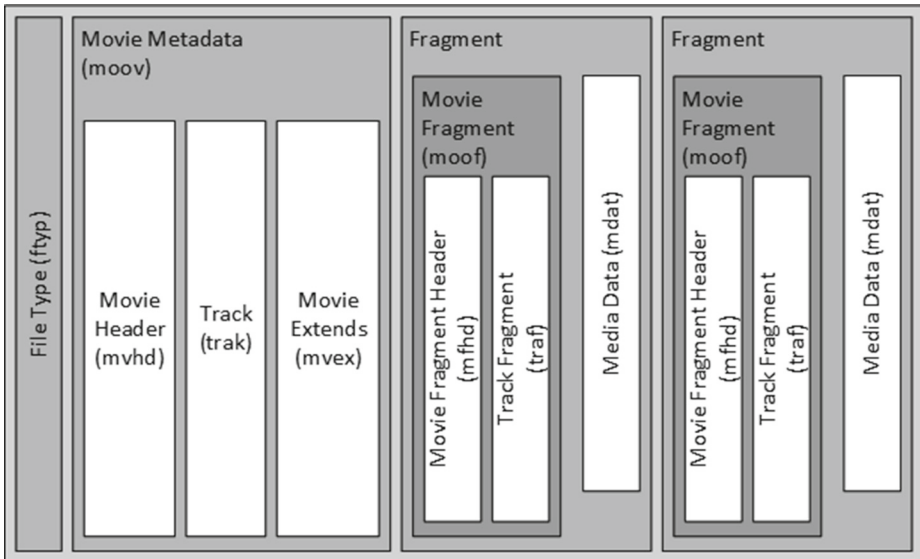


Fig. 1. MP4 fragment format

The fragmented MP4 is currently used for YouTube’s range request algorithm and is currently the default format for newly uploaded videos. Therefore the application used for the YouTube buffer prediction focuses on this format.

To detect the requests to the server, the outgoing TCP traffic is monitored for HTTP GET requests. When a request is detected, video information like the video id, the format and the signature as well as streaming related information like the utilized flow control approach, are derived.

As for accumulating the currently downloaded amount of data for a certain video, the incoming traffic is monitored for the fragment boxes and incoming fragments are added to the estimation of the video’s download process. The buffer estimation is then calculated as the difference between the available playback time of the already downloaded amount and the time passed since the begin of playback. We compute the application state for the video and the audio flow. User interaction like pausing, reverse jumping can not be taken into account. Forward jumping can be identified using the media data if the user jumps to a location which has not been downloaded yet.

3.2 Enforcing Resource Control

The most flexible way to control the network resources on a per-flow basis is the assignment of a dynamic rate limit to the individual flows. This enables a granular adjustment of each flow, but also results in a lot of state information required to maintain the single queues for each individual flow. A more simple approach is to define several priority queues and assign the flows to the different priorities with respect to the current application state, cf. Fig. 2. Since we evaluate a small home scenario, we implemented the priority approach by using five different queues with

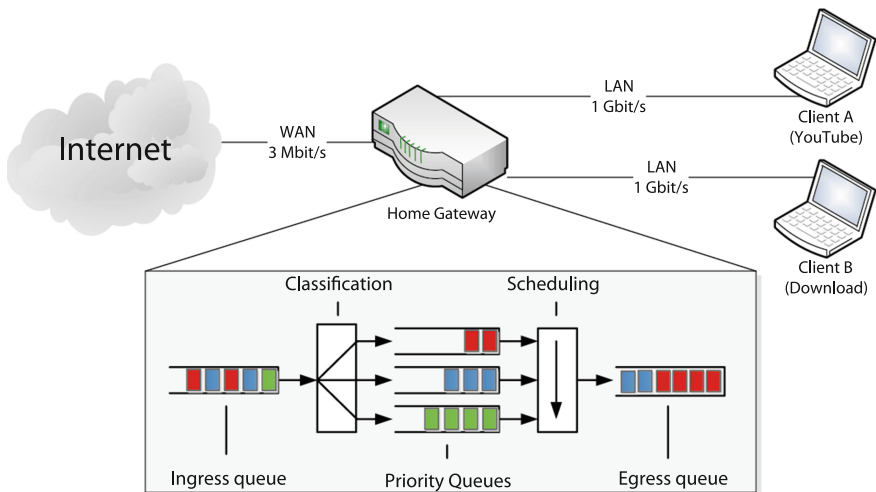


Fig. 2. Buffered playtime for audio and video

a queue size of 25 packets for each queue. To dynamically re-assign flows to different priorities, we use a Python wrapper script for the Linux Traffic Control [10] (TC) API via a simple TCP socket interface. With TC a stateful queuing is done that classifies the incoming packets and sorts them into the priority classes that are emptied by a scheduler according to the priorities.

3.3 Triggering Resource Control Actions

The network controller is implemented within the application monitor and executed periodically. If the YouTube video or audio playback buffer falls below 25 s, it is moved to a higher priority queue. If they overrun 40 s, the flow is moved to a lower priority. Thus, a reaction takes place if the application quality is endangered, i.e., the playback buffers threaten to run empty. Best-effort behavior is used, as long as the application quality for all applications is not endangered.

3.4 Application Design

The described components are implemented in two binaries. One is the resource enforcement application using priority queues. The other components are implemented as a C++ linux application. The multi-threaded design of the application together with the compiler optimizations provided by C++ allow low latencies and a high performance on the used end device. We use the local TCP/IP interface to perform the communication between the applications. The redirection of the YouTube flow through the user space, however, has a significant impact on the data plane performance, since the CPU of the home gateway becomes the bottleneck. This results in a CPU load of 40 % for a link capacity of 3 MBit.

4 Experimental Setup and Measurement Procedure

In the following subsection, the measurement scenario and the measurement sequence is described in order to evaluate the resource control.

4.1 Experimental Setup

The implemented application is designed to run on a typical home gateway. Among the available off-the-self gateway hardware, an AVM Fritz!Box 3390 has been chosen, featuring a MIPS 34Kc CPU running at 300 MHz with 22 MiB of RAM. To allow the usage of custom software and to provide the default linux libraries the alternative open source operating system Freetz [11] is used. On the Fritz!Box the resource control application and the combined decision and monitoring application are started.

The gateway is connected to the Internet via Ethernet with a theoretical bandwidth of 100 Mbit/s. In order to emulate a limited connection to the Internet, we restrict the downlink from the Internet to 3 Mbit/s. Two client laptops (Client A and Client B) are connected to the gateway via Ethernet with a link speed of 1 Gbit/s.

4.2 Measurement Procedure

For our investigations, we use the following procedure. The scenario begins with starting a YouTube video clip on Client A. After 15 s a download is started on Client B utilizing four parallel connections for the download. Hence, we can expect the download to get at least four times the bandwidth of the YouTube stream. As video clip, we always use the same video clip, “VAN CANTO - Badaboom (Official)”¹ in 720p. After 300 s we stop the measurement.

5 Evaluation Results

5.1 YouTube Buffer Level During Video Playback

At first, we have a look at the YouTube buffer level over time without any resource management. Using fragmented MP4, the audio and the video stream are transmitted separately resulting in different playback buffer states. If one of the playback buffers is empty, the video clip stalls until the stalling threshold of 5 s is reached.

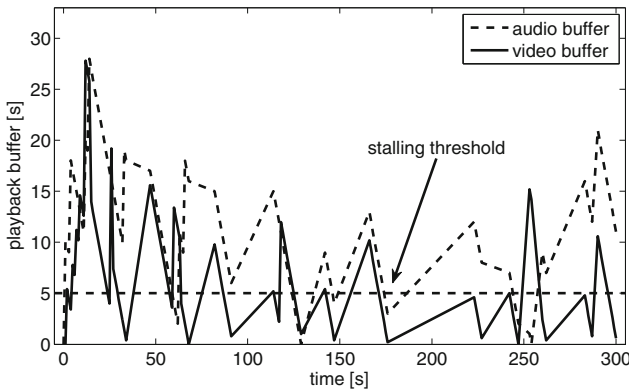


Fig. 3. Buffered playtime for audio and video without application monitoring and resource management

Figure 3 shows the playback buffer for the audio and video buffers. Both buffer levels increase until the parallel download is started. After that, the buffer level for both buffers decreases at about 20 s. In this example, the video buffer initially is empty at second 40, resulting in a short stalling period although the audio buffer stays at a sufficient buffer level. Since not enough network capacity is available to allow a smooth video playback in case of a concurrent file download, the video playback is interrupted several times. After 60 s, the audio buffer falls

¹ Video clip available at <http://www.youtube.com/watch?v=Aeaz4s7q0Ag&wide=1&hd=1>.

below the threshold of 5 s while the video buffer still maintains a higher buffer level. Consequently, the audio flow is also taken into account in the following.

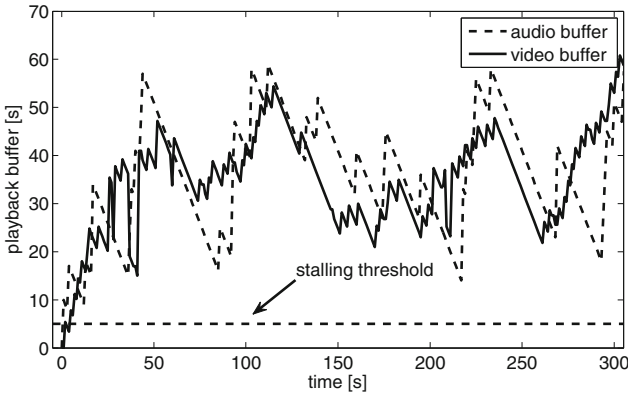


Fig. 4. Buffered playtime for audio and video with application monitoring and resource management

Figure 4 shows audio and video playback buffer over time with dynamic prioritization enabled. In contrast to the best-effort case, the buffer levels increase although the parallel download has started. Both flows are prioritized until the playback buffers overrun a threshold of 45 s. After that, the resource control mechanism is turned off, until the buffers fall below 25 s. The TCP flow control results in additional dynamics and it may take some time until a fair bandwidth share between the flows is reached. This is reflected by the minimal and maximal playback buffer levels at 15 s and 63 s. In addition, the re-assignment to different priorities may lead to packet reordering influencing the TCP control loop and resulting in a reduction of the TCP sending window, cf. [12].

5.2 Statistical Evaluation of the Investigated Scenario

After investigating the time series for the audio and video streams for the scenarios with and without prioritization, we focus on a statistical significant comparison between both approaches. For that, we conduct 10 runs and compare the CDFs for the buffer fillings. The results of this investigation is illustrated in Fig. 5 with a confidence level of 95 %. It can be seen that the application-aware approach using dynamic prioritization with respect to the application state clearly outperforms the approach without prioritization, i.e., the best effort case. This holds for both flows, the audio and the video flow. Stalling is minimized allowing a better user-perceived quality for the video streaming user. Further, it can be seen, that the buffer level for the audio buffer is typically higher as for the video buffer. Hence, we can conclude that a video stalling is more likely due to a video

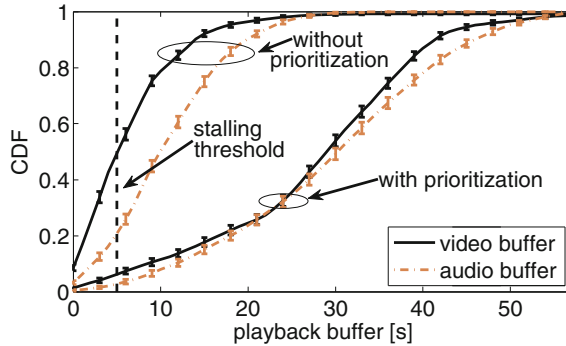


Fig. 5. Buffered playtime for audio and video

buffer under run. Figure 5 also shows that the majority of values are located between 20 and 50 s, closely related to the values of the control hysteresis.

6 Related Work

In the following, papers on different resource management approaches are summarized. We specifically address QoE cross-layer resource management and application-aware resource management.

Different common techniques for resource management are used in networking devices and current communication protocols [13–15]. They all make use of protocol-specific conditions to favor packets or IP flows, or to influence the data transmission. To identify relevant packets or streams other techniques are necessary. Approaches like deep packet inspection (DPI) or explicit signaling of the application [16] are used to identify different applications and to map them to QoS classes. Approaches like [17, 18] tag the ToS (type of service), the diffserv (differentiated services) field in the IP packet or use a shim header [16] to prioritize the corresponding network traffic.

QoE Cross-Layer Resource Management. The goal of QoE resource management approaches is a resource allocation according to the QoE of the users. QoE models [19, 20] for different applications form the basis as proposed in [21–25]. In order to incorporate QoE, a cross-layer optimization is typically used [21, 22, 25]. For example, in case of overload on the network layer, [21, 22] prioritize important packets of MPEG videos, mainly I-frames, in order to still guarantee a high quality. In [26], a multi-layer video encoding with scalable video codec is used. The goal is to prioritize the different layers in different QoS classes to specifically drop video layers with less importance if the network is congested. In [23], QoE-based scheduling for wireless mesh networks is proposed. According to a MOS metric that maps QoE to QoS parameters, the video, data, and audio traffic is forwarded.

Application-aware resource management. In addition to the work mentioned above, application-aware resource management is seen as a step towards QoE-oriented or QoE-aware resource management. It allocates resources based on dynamic cross-layer information from the application, e.g., the buffered play-time in case of videos in order to increase QoE [27–29]. This approach can be seen as a continuation of former work which used a so-called utility function for network resource management to define a scheduling order with respect to different application cases [30–34].

In [35], a QoS manager is proposed which centrally passes the requirements for the whole network to a switch. The problem is addressed how various networking entities can work together to prioritize traffic flows in order to allow a prioritization for the entire network. For residential networks, [36] proposes to setup a QoS system that uses hints from individual hosts to make decisions about traffic prioritization. In [37], the concept for traffic prioritization is presented for small networks. The authors specify a prioritization of applications for active users and a de-prioritization for applications in the background.

7 Conclusion

Traffic in today’s network structures is typically 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. This may result in unfairness in terms of the user-perceived quality, in particular if the overall resources are limited. This holds especially for home networks where different users may compete for limited network resources.

In this paper, we presented an application-aware networking approach to overcome this problem. We investigated the feasibility of the approach for a scenario consisting of two users, a download user and a user watching a YouTube video clip. We implemented an application monitoring component for YouTube, a prioritization mechanism to control the resources, and a simple decision logic. The monitoring component is able to estimate the video and audio playback buffer. The resulting information are used to trigger a prioritization of the video streaming if the buffer runs empty. Hence, a good QoE for the video streaming user can be guaranteed. The components were implemented on a typical home gateway. The evaluation of the scenario indicates the potential of the mechanism to manage the application quality and therewith the QoE for multiple users in a multi-application scenario.

References

1. Home Gateway Initiative: Home Gateway QoS module requirements. whitepaper, Dec 2012
2. Eckert, M., Knoll, T.M.: ISAAR (Internet Service Quality Assessment and Automatic Reaction) a QoE monitoring and enforcement framework for internet services in mobile networks. In: Timm-Giel, A., Strassner, J., Agüero, R., Sargento, S., Pentikousis, K. (eds.) MONAMI 2012. LNICST, vol. 58, pp. 57–70. Springer, Heidelberg (2013)

3. Wamser, F., Zinner, T., Iffländer, L., Tran-Gia, P.: Demonstrating the prospects of dynamic application-aware networking in a home environment. In: Proceedings of the 2014 ACM Conference on SIGCOMM, pp. 149–150. ACM (2014)
4. Alcock, S., Nelson, R.: Application flow control in youtube video streams. *ACM SIGCOMM Comput. Commun. Rev.* **41**(2), 24–30 (2011)
5. Sieber, C., Hofffeld, T., Zinner, T., Tran-Gia, P., Timmerer, C.: Implementation and user-centric comparison of a novel adaptation logic for DASH with SVC. In: IFIP/IEEE International Workshop on Quality of Experience Centric Management (QCMAN), Ghent, Belgium (2013)
6. Wikipedia: Youtube – Wikipedia, the free encyclopedia (2014). <https://en.wikipedia.org/wiki/YouTube>. Accessed 20 May 2014
7. A.S. Incorporated: Video File Format Specification, Adobe Systems Incorporated Std., Rev. 10, November 2008. http://download.macromedia.com/f4v/video_file_format_spec_v10_1.pdf. Accessed 26 May 2014
8. Singer, D.: 3GPP TS 26.244; Transparent end-to-end packet switched streaming service (PSS); 3GPP file format (3GP), ETSI 3GPP Std., Rev. 12.3.0, March 2014. <http://www.3gpp.org/DynaReport/26244.htm>. Accessed 26 May 2014
9. MPEG 4 standards ISO/IEC 14496–1 ff, International Organization for Standardization Std., Rev. 2010 (1999). http://www.iso.org/iso/iso_catalogue/catalogue_ics/catalogue_detail_ics.htm?csnumber=24462
10. Hubert, B.: Linux Advanced Routing & Traffic Control, Linux Foundation. Accessed 26 May 2014
11. Freetz project: Freetz. <http://freetz.org/>
12. Zinner, T., Jarschel, M., Blenk, A., Wamser, F., Kellerer, W.: Dynamic application-aware resource management using software-defined networking: implementation prospects and challenges. In: IFIP/IEEE International Workshop on Quality of Experience Centric Management (QCMAN), Krakow, Poland (2014)
13. IEEE 802.1Q 2011: Standard for Local and Metropolitan Area Networks - Media Access Control (MAC) Bridges and Virtual Bridge Local Area Networks, August 2011
14. IEEE 802.11e-2005: Standard for Information technology - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications, November 2005
15. IEEE 802.16m-2011: Standard for Local and metropolitan area networks - Part 16: Air Interface for Broadband Wireless Access Systems, Amendment 3: Advanced Air Interface (802.16m-2011), May 2011
16. Paul, S., Jain, R., Pan, J., Iyer, J., Oran, D.: Openadn: A case for open application delivery networking. In: 2013 22nd International Conference on Computer Communications and Networks (ICCCN), pp. 1–7. IEEE (2013)
17. RFC 791: DARPA Internet program protocol specification (1981)
18. Nichols, K., Blake, S., Baker, F., Black, D.: RFC 2474: Definition of the differentiated services field (DS field) in the IPv4 and IPv6 headers (1998)
19. Agboma, F., Liotta, A.: QoE-aware QoS management. In: 6th International Conference on Advances in Mobile Computing and Multimedia, pp. 111–116. ACM (2008)
20. Fiedler, M., Hofffeld, T., Tran-Gia, P.: A generic quantitative relationship between quality of experience and quality of service. *IEEE Network* **24**(2), 36–41 (2010). Special Issue on Improving QoE for Network Services
21. Gross, J., Klaue, J., Karl, H., Wolisz, A.: Cross-layer optimization of OFDM transmission systems for MPEG-4 video streaming. *Comput. Commun.* **27**(11), 1044–1055 (2004)

22. Khan, S., Peng, Y., Steinbach, E., Sgroi, M., Kellerer, W.: Application-driven cross-layer optimization for video streaming over wireless networks. *IEEE Commun. Mag.* **44**(1), 122–130 (2006)
23. Reis, A., Chakareski, J., Kessler, A., Sargento, S.: Quality of experience optimized scheduling in multi-service wireless mesh networks. In: *IEEE Conference on Image Processing (ICIP)*, pp. 3233–3236. IEEE (2010)
24. Pries, R., Hock, D., Staehle, D.: QoE based bandwidth management supporting real time flows in IEEE 802.11 mesh networks. *Praxis der Informationsverarbeitung und Kommunikation* **32**(4), 235–241 (2010)
25. Ameigeiras, P., Ramos-Munoz, J.J., Navarro-Ortiz, J., Mogensen, P., Lopez-Soler, J.M.: QoE oriented cross-layer design of a resource allocation algorithm in beyond 3G systems. *Comput. Commun.* **339**(5), 571–582 (2010)
26. Huang, C., Juan, H., Lin, M., Chang, C.: Radio resource management of heterogeneous services in mobile WiMAX systems [Radio Resource Management and Protocol Engineering for IEEE 802.16]. *IEEE Wireless Commun.* **14**(1), 20–26 (2007)
27. Thakolsri, S., Khan, S., Steinbach, E., Kellerer, W.: QoE-driven cross-layer optimization for high speed downlink packet access. *J. Commun.* **4**(9), 669–680 (2009)
28. Superiore, L., Wrulich, M., Svoboda, P., Rupp, M., Fabini, J., Karner, W., Steinbauer, M.: Content-aware scheduling for video streaming over HSDPA networks. In: *Second International Workshop on Cross Layer Design, IWCLD 2009*, pp. 1–5. IEEE (2009)
29. Staehle, B., Wamser, F., Hirth, M., Stezenbach, D., Staehle, D.: AquareYoum: Application and quality of experience-aware resource management for YouTube in wireless mesh networks. *PIK - Praxis der Informationsverarbeitung und Kommunikation* (2011)
30. Xiao, M., Shroff, N., Chong, E.: A utility-based power-control scheme in wireless cellular systems. *IEEE/ACM Trans. Netw.* **11**(2), 210–221 (2003)
31. Andrews, M., Qian, L., Stolyar, A.: Optimal utility based multi-user throughput allocation subject to throughput constraints. In: *IEEE INFOCOM*, vol. 4, pp. 2415–2424. IEEE (2005)
32. Song, G., Li, Y.: Utility-based resource allocation and scheduling in ofdm-based wireless broadband networks. *IEEE Commun. Mag.* **43**(12), 127–134 (2005)
33. Saul, A.: Simple optimization algorithm for mos-based resource assignment. In: *VTC Spring 2008*, pp. 1766–1770. IEEE (2008)
34. Pei, X., Zhu, G., Wang, Q., Qu, D., Liu, J.: Economic model-based radio resource management with qos guarantees in the cdma uplink. *Eur. Trans. Telecommun.* **21**(2), 178–186 (2010)
35. Zinca, D., Dobrota, V., Vancea, C., Lazar, G.: Protocols for communication between qos agents: Cops and sdp. In: *COST*
36. Katchabaw, M., Lutfiyya, H., Bauer, M.: Usage based service differentiation for end-to-end quality of service management. *Comput. Commun.* **28**(18), 2146–2159 (2005)
37. Martin, J., Feamster, N.: User-driven dynamic traffic prioritization for home networks. In: *Proceedings of the 2012 ACM SIGCOMM Workshop on Measurements Up the Stack*, pp. 19–24. ACM (2012)