

Modelling and performance analysis of application-aware resource management

Florian Wamser, Andreas Blenk, Michael Seufert, Thomas Zinner, Wolfgang Kellerer, Phuoc Tran-Gia

Angaben zur Veröffentlichung / Publication details:

Wamser, Florian, Andreas Blenk, Michael Seufert, Thomas Zinner, Wolfgang Kellerer, and Phuoc Tran-Gia. 2015. "Modelling and performance analysis of application-aware resource management." *International Journal of Network Management* 25 (4): 223–41.
<https://doi.org/10.1002/nem.1894>.

Nutzungsbedingungen / Terms of use:

licgercopyright

Dieses Dokument wird unter folgenden Bedingungen zur Verfügung gestellt: / This document is made available under these conditions:

Deutsches Urheberrecht

Weitere Informationen finden Sie unter: / For more information see:

<https://www.uni-augsburg.de/de/organisation/bibliothek/publizieren-zitieren-archivieren/publiz/>



Modelling and performance analysis of application-aware resource management

Florian Wamser,^{1,*†} Andreas Blenk,² Michael Seufert,¹ Thomas Zinner,¹
Wolfgang Kellerer² and Phuoc Tran-Gia¹

¹*Institute of Computer Science, University of Würzburg, Würzburg, Germany*

²*Department of Electrical and Computer Engineering, Technische Universität München, Munich, Germany*

1. INTRODUCTION

The way how customers use today's communication networks has changed dramatically in recent decades. In the past, users ran elastic and less interactive applications, such as file transmissions or emails. Today, they run a huge set of high-quality applications, which are dominating today's network traffic, for example, video streaming, online games and social networking. In particular, for high-quality applications, users do not require high data rates; they require a good application quality of experience (QoE). Among other subjective and objective factors that influence QoE (user expectations, hardware device capabilities, etc.), the applications have different and varying demands on the network in terms of QoE. This includes, for example, low latency for interaction, high data rate for data retrieval or low packet loss for live content streaming. Consequently, just optimizing network parameters does not necessarily lead to QoE improvement without considering applications needs.

*Correspondence to: Florian Wamser, Institute of Computer Science, University of Würzburg, Würzburg, Germany.

†E-mail: florian.wamser@informatik.uni-wuerzburg.de

One solution in order to meet the application requirements is to employ over-provisioning in the network. As pure over-provisioning of network resources typically results in high costs and thus is not efficient from the economical point of view, new concepts have to be considered in order to improve networks' resource efficiency and application QoE simultaneously. As network operators do not have control over the server or customer premises, they have to introduce new concepts inside their networks. Application-aware resource management is such a concept that has emerged as a first step towards improving network resource efficiency in terms of QoE [1–9]. The idea of application-aware resource management is to orchestrate networks according to user demands that are based on application quality indicators with high impact on users' QoE. Network resource management and orchestration actions are, for example, changing the routing of application flows among the network or configuring the resource allocation of particular network links, for example, by using different resource scheduling such as priority queuing or weighted fair queuing. In this way, the utility per allocated network resources in terms of QoE can directly be influenced. The quality indicators can be fetched from the network based on in-network measurements [10,11]. A typical implementation of the concept consists of three entities: (i) the network and application monitoring deriving the key performance indicators and network characteristics [9,12]; (ii) a decision unit [9]; and (iii) a resource management action that changes the traffic handling within the network (e.g. traffic scheduling with queuing disciplines, routing and dynamic path selection) or the traffic produced by the applications (e.g. lowering video resolution/quality and server selection) [5,9,13]. Although network resources may be shared unequally, such a traffic handling even leads to an improved overall performance for all users in terms of QoE [2].

In existing work [2,14,15], the benefit of application-aware resource management mechanisms in terms of network resource efficiency and user satisfaction was shown for certain applications. However, the trade-offs between different application-aware resource management mechanisms for network links in a multi-application scenario in terms of complexity and implementation level have not been investigated. In detail, most research has been use-case driven and evaluated solutions for only tiny scenarios. The existing concepts have not provided a holistic evaluation approach that tries to understand underlying concepts of how to technically realize resource allocation algorithms for a multi-application scenario under varying traffic conditions. An algorithm designed for a particular setup may not be efficient under all circumstances. Consequently, a systematic analysis of application-aware resource management is important to quantify the real gain of the concept, in particular, when comparing with alternative solutions, such as network over-provisioning.

This paper provides methods that allow network operators quantifying the benefit of application-aware resource allocation algorithms for network links in three dimensions: for web browsing and progressive video streaming applications for constant bit rate and variable bit rate streams; for varying traffic situations and mixes; and for four algorithms covering the space of application-aware resource allocation. This can be carried out by the proposed analysis algorithms and the simulation, by inserting the own network settings and traffic situations. Although an application-aware resource management should consider different network operations, such as changing the routing or server selection, we focus on resource allocation for network links, as there is already a high potential for improvement. In addition, improvements can be simply implemented by the operator, because only one link is touched, which is under the responsibility of the operator. The implementation complexity varies in terms of actions during run time and the needed information about the applications. As we are interested in improving the performance for one network link, we abstract the link by a queuing model. By using the model, application parameters can be derived, which allow to draw conclusions about the resulting application QoE. The prerequisite is that this link represents a bottleneck in the network. In particular, only problems on the bottleneck link are addressed in this paper. Other influences outside the bottleneck link cannot be solved and are not in focus of the algorithms considered in this paper. To conduct the evaluations, analytic models are used that are based on approaches using a Markov M/M/1 queue with processor sharing policies (M/M/1-PS) [16,17]. In case of algorithms designed for managing dynamic applications like progressive video streaming with variable bit rate, a discrete-event simulator implemented in Java is used. Using both analytic and simulative approaches, a study of all algorithms for different traffic scenarios is performed in order to give selection and configuration guidelines for the algorithms.

In particular, the paper highlights the following contributions:

1. modelling the application-aware resource management problem;
2. modelling four application-aware resource allocation algorithms with different complexity;
3. analytical and simulative analysis of the algorithms;
4. guidelines for the selection and configuration of the algorithms.

The structure of the paper is as follows. First, we outline related work and discuss resource management options for operators in access networks in Section 2. We introduce the model for the considered scenario and formulate the resource allocation problem in Section 3. In Section 4, we introduce four resource allocation algorithms based on the application-aware concept. In Section 5, we explain the analytic and simulative methodologies for the evaluation of the algorithms. In Section 6, we show the results of performance evaluation for the algorithms in terms of network load, traffic mix and the aforementioned quality indicators. Finally, we draw conclusions in Section 7.

2. BACKGROUND AND RELATED WORK

Today, the dimensioning of the capacity of access networks is a crucial task for operators. On the one hand, the cost must be kept as low as possible (*cost-efficiency*). On the other hand, sufficient capacity must be available for any expected application or service to satisfy the users (*QoE improvement*). Although it was widely believed that network traffic may not be predictable, currently published measurements in literature about access networks show a predictable network traffic behaviour [18,19]. Based on such a predictable behaviour, an operator can analyse the trade-offs before upgrading the network. According to the traffic mix and patterns in the network, the best setup can be chosen in order to improve the overall network performance in terms of user-perceived service quality. One possibility to achieve this is the methodology presented in this work for *application-aware resource management*.

In this section, we summarize related work, which is subject to application-aware resource management. The current state of the art consists of numerous approaches and concepts, which use information from different networking layers and of different degrees of complexity in terms of their functionality, the information flow and the number of involved entities.

Interfaces and application programming interfaces for application-aware resource management.

The currently prevalent idea in resource management for improving the quality for end-users is to differentiate traffic flows into quality of service levels. In [20], an application network interface is proposed for OpenFlow-based networks to specify different service levels in the network. The interface allows application service providers to define application-specific requirements that are implemented by the network. To achieve a network-wide enforcement, Application Label Switching is introduced that adds additional protocol headers (*labels*) to IP flows that can be utilized by OpenFlow software-defined networking switches in order to differentiate traffic at bottleneck links. The advantage of this approach is the usage of the standardized OpenFlow protocol that defines existing bandwidth provisioning algorithms such as weighted deficit round-robin or weighted fair queuing for resource management. However, a comprehensive quantification of the benefits with respect to the different algorithms is lacking. Similar to this approach, an application network interface via an additional protocol layer is proposed in [21]. This service access layer facilitates a split between service level control and the application data plane, thus enabling service-aware network routing. For the realization, the proposed protocol must be implemented on each network device. In [22], 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. Finally, there are specific QoS-based or QoE-based management frameworks that define network-wide elements to coordinate the traffic forwarding. For residential networks, Katchabaw *et al.* [23] propose to set up a QoS system that uses hints from individual hosts to make decisions about traffic prioritization.

Application-aware resource management with central entity. Besides [23], the authors introduce in [24] a central entity that has knowledge about the network and application situation. It performs a dynamic resource allocation decision that utilizes software-defined networking in order to enforce

application demands in the network. An implementation featuring a video on demand application competing with a file download at a bottleneck link shows that monitoring the video on demand application's buffer state can help improving the end-users' QoE. The work on participatory networking presented in [1] proposes a communication mechanism between applications and the network, which is initiated by the applications. Key concepts include conflict resolution between the needs of different applications or users and decomposition of network control, that is, limiting an application's authority. The authors of [2] introduce an OpenFlow-based framework that aims at achieving fairness with respect to QoE among all users in an adaptive video streaming context. The expected QoE of a user is estimated on the quality of information during the video streaming. This approach combines a network-aware service with a network-based control entity. Based on monitoring results gathered via intercepting media presentation description files used in Dynamic Adaptive Streaming over HTTP (DASH), the framework dictates video quality levels for each user in order to maximize the resulting QoE for every participant. Despite the achieved fairness, open issues here are how universal this approach is and how it can be applied to other applications.

A system design featuring functional blocks that represent functions for packet handling in the network is presented in [9,25]. Feedback loops including a central decision unit allow for combining network-aware applications with an application-aware network. A central idea consists of translating user demands into application demands, which in turn can be translated into network requirements. The system's capabilities are demonstrated and verified for various use cases including live video and video on demand scenarios.

Application-aware resource management for video streaming. In [3], the authors propose to consider the buffered playtime of a video when assigning network resources. They propose algorithms optimizing the delivered video quality. In [2], the authors propose to maximize the QoE of the end-users running adaptive video streaming applications. In [8], the authors propose to adapt the video rate based on client-side feedback of the buffered video time.

Multi-application optimization. Huang *et al.* [26] consider video streaming, voice over IP and data services when optimizing radio resource management for mobile WiMAX systems. However, they consider only the requested bandwidth per application and do not differentiate between different application quality levels. In [27], the authors focus on cross-layer optimization between application layer and medium access control layer. The application parameters are given by a generic mathematical model for the Mean Opinion Score (MOS) as a metric for the perceived quality. Reis *et al.* [28] consider video streaming, audio streaming and data services in a wireless mesh network. The authors' proposed multidimensional optimization tries to minimize the distortion across flows on wireless links under given fairness constraints. In [29], the authors propose resource allocation for Long-Term Evolution based on a particle swarm optimization while considering multiple applications.

Application-Layer Traffic Optimization. A protocol realizing the exchange of information between network and application is standardized by the Internet Engineering Task Force Application-Layer Traffic Optimization (ALTO) working group [30]. It aims at providing guidance to content delivery applications such as peer-to-peer or content delivery networks, which have to select one or several hosts or endpoints out of a set of possible candidates. Such appropriate candidates can be selected, and the performance with respect to user-centric, network-centric and application-centric metrics can be improved. Currently, several extensions of the ALTO protocol including data centres and cloud applications are being discussed [31].

Data centre applications. More work on application-aware resource management has been conducted in the context of data centre applications. Research on data centre architectures already shows that combined solutions, for example, for traditionally separated mechanisms, such as routing and service migration, may increase data centre efficiency. However, data centre applications are not end-user applications as in access networks but run in isolated and manageable server farms. In general,

the constrained environment of data centres enables the combination of these traditionally separated mechanisms. Data centre application-aware solutions are also provided in [32–34].

In summary, there are many approaches in the literature that propose user-aware or application-aware resource management. All related works, however, do not evaluate the impact of their algorithms in a global fashion with different network loads or different application traffic mixes, that is, a variable proportion of different applications on the network use. The reason for this is simple: modelling of the problem and a theoretical abstraction are necessary, besides the practical use-case driven evaluations, to be able to perform more extensive studies.

3. MODEL

In this section, we begin by describing the system of interest—a bottleneck link shared by multiple users running either browsing or video streaming applications. Following the model, we define a resource allocation objective for the network scenario. The symbols, which are used throughout this work, are summarized in Table 1.

3.1. Network and applications

A bottleneck link l has a limited capacity C . While sharing the link, users are running either web browsing applications, which we identify via w , or video streaming applications, which we identify via v . Each application requests content of a certain size. The size of each application request is given by s_w and s_v , respectively. Furthermore, a video also has a specific video length t_v . All active web flows and all active video flows build the set of active application flows $\mathcal{A} = \{a_1, \dots, a_N\}$. Each flow may occupy a certain amount of the link's capacity. The request times of the users for web w and for video streaming v are modelled as Poisson processes, that is, they arrive with rates λ_w and λ_v .

To assess the performance of the applications in the given network, we evaluate specific key performance indicators for each application [35,36]. In the case of web, the page load time of a web site has the most dominant impact on perceived QoE [35]. Therefore, we consider the download time of a web page as the key performance indicator of web browsing. For video streaming, there are several well-known metrics and evaluation criteria: Structural SIMilarity (SSIM), Peak Signal-to-Noise-Ratio (PSNR), number of stallings during playback and length of stalling periods. We investigate hyper-text transfer protocol (HTTP) progressive video streaming in this paper. Objective metrics such as SSIM and PSNR consider only the quality of each single video frame. This is valid and useful for live streaming without buffering periods where image distortion can occur. For HTTP progressive video streaming, the image quality itself is never impaired because of the use of the TCP. If the bandwidth

Table 1. Summary of used symbols for resource allocation problem formulation, resource allocation algorithms and analytic performance evaluation model

Symbol	Definition
\mathcal{A}	The set of active application flows
$a \in \{v, w\}$	An active video or web flow
$S \subset \mathcal{A}$	A subset of active flows
S_v, S_w, S_p	The set of active video/web/prioritized flows
C	Link capacity
q	Allocated bandwidth
q_a, q_S, q_p	Bandwidth allocated to flow/set of flows/set of prioritized flows
\bar{b}	Average video bit rate
σ_b	Standard deviation of video bit rate
s	Flow size
λ	Arrival rate
μ	Service rate
ρ	Utilization

is not sufficient, the video playback simply stops (stalling) and the picture freezes until the buffer is refilled. According to [36,37], for HTTP progressive video streaming, stalling (i.e. playback interruptions) is the most dominant factor of QoE, clearly exceeding the significance of video resolution as a second impact factor. Thus, we investigate the buffering ratio, that is, the total stalling time t_s divided by the total video length t_v , as a performance indicator of video flows.

3.2. Resource allocation

TCP is the dominant transport protocol for both application types. Because of TCP, all active application flows $\mathcal{A} = \{a_1, \dots, a_N\}$ may share the link capacity C equally. We assume that the share ratio can additionally be influenced by a network algorithm, that is, a resource allocation policy, for each flow $a_i \in \mathcal{A}$ or for each disjoint set of flows $S_j \subset \mathcal{A}$. This means that there is the possibility of assigning each flow a_i a bandwidth q_{a_i} with $\sum_{i=1}^N q_{a_i} \leq C$. Similarly, each set of flows S_j can be assigned a bandwidth q_{S_j} , such that $\sum_{S_j \subset \mathcal{A}} q_{S_j} \leq C$. The bandwidth of a set of flows can be again shared equally such that each flow $a_j \in S_j$ has bandwidth $q_{a_j} = \frac{q_{S_j}}{|S_j|}$.

Different application characteristics, usage behaviours and QoE factors have to be considered when assigning bandwidths to the active flows. Thus, the problem on which we focus can be formulated as follows: Find a mapping $\mathcal{A} \rightarrow \mathbb{R}, a_i \rightarrow q_{a_i}$, such that $\sum_{i=1}^N q_{a_i} \leq C$ holds, and the average QoE of all active flows a_i , $i \in \{1, 2, \dots, N\}$ is maximized and fair in terms of the objective given by the application type.

As described in Section 3.1, the key performance indicators of each application will be evaluated that have a high correlation to the QoE. This means that a well-performing resource allocation algorithm is supposed to achieve short download times for web flows and low buffering ratios for video streams.

4. RESOURCE ALLOCATION ALGORITHMS

This section introduces four algorithms aiming at an improved resource allocation in terms of perceived QoE of the end-user. In the following, we take only two applications, that is, video streaming and web browsing, into account.

4.1. Fixed bandwidth allocation for all video flows (FBV)

The FBV algorithm reserves a fixed bandwidth for video traffic as it is the more demanding application. The reserved capacity q_{S_v} is a parameter that can be adjusted to reflect the actual traffic share of video traffic.

Algorithm: All currently active flows are split into two sets, one set S_v containing all video flows and one set S_w containing all web browsing flows. S_v is allocated a certain bandwidth q_{S_v} , which is shared equally among all active video flows. This means that each video flow obtains a bandwidth of $q_v = \frac{q_{S_v}}{|S_v|}$. Consequently, S_w is assigned $q_{S_w} = C - q_{S_v}$, which is also shared equally, such that each web browsing flow receives $q_w = \frac{q_{S_w}}{|S_w|}$. In order to avoid starvation of web browsing, q_{S_v} has to be smaller than C . If one of the two sets S_v or S_w is empty, the other set can share the whole link capacity C among its flows. This means that $q_v = \frac{C}{|S_v|}$ if S_w is empty, and $q_w = \frac{C}{|S_w|}$ if S_v is empty.

Figure 1(a) illustrates the resource allocation of FBV for two video flows (v_1 and v_2) and one web flow w_1 . The video flows are sharing q_{S_v} equally, and the web flow receives the remaining capacity $C - q_{S_v}$. In case one set S_v or S_w is empty, the applications share C equally, as illustrated in Figure 1(b).

4.2. Fixed bandwidth allocation for each video flow (FBF)

A fixed bandwidth can be reserved not only for the entire set of video flows (FBV) but also for each individual flow. Thus, specific characteristics of each flow can be taken into account. In the performance evaluation in Section 6, the reserved capacity q_{v_i} for a video flow a_i will depend on video




Figure 1. Behaviour of FBV for two different cases in terms of number of application flows and types of application flows




Figure 2. Behaviour of FBF for two different cases in terms of varying requested data rate by the video flows

attributes such as its video bit rate b_i . For example, for videos having CBR, the average video bit rate ($q_{v_i} = \bar{b}_i$) is allocated. In the case of videos that show a variable bit rate (VBR) behaviour, the average video bit rate plus its standard deviation ($q_{v_i} = \bar{b}_i + \sigma_{b_i}$) is assigned to additionally account for bit rate variations of VBR content. However, different characteristics and bandwidth allocations could also be used in this algorithm. Technically, the FBF algorithm also has to prevent starvation of web browsing by reserving a minimum bandwidth q_{S_w} (parameter) and adjust the allocated bandwidths in case more bandwidth than available is requested.

Algorithm: All currently active flows are split into two sets, one set S_v containing all video flows and one set S_w containing all web browsing flows. Each active video flow $a_i \in S_v$ can request a certain bandwidth q_{v_i} . If the remaining available capacity $q_r = C - q_{S_w} - \sum_{a_j \in S_v} q_{v_j} \geq 0$ (i.e., all requests can be fulfilled), the remaining capacity q_r is added to q_{S_w} . Each video flow a_i is assigned its requested bandwidth q_{v_i} , and all web browsing flows equally share $q_{S_w} + q_r$. Thus, each web browsing flow is assigned $q_w = \frac{q_{S_w} + q_r}{|S_w|}$. In case $q_r < 0$, more bandwidth is requested than available. Therefore, all web flows equally share only their guaranteed minimum bandwidth ($q_w = \frac{q_{S_w}}{|S_w|}$), and the capacity $C - q_{S_w}$ is assigned to the video flows relatively to the requested bandwidths. Thus, each video flow a_i is allocated a bandwidth of $q'_{v_i} = (C - q_{S_w}) \cdot \frac{q_{v_i}}{\sum_{a_j \in S_v} q_{v_j}}$.

Figure 2(a) shows the bandwidth allocation in case that the requested video rates $q_{v_1} + q_{v_2} \leq q_{S_v}$. The video flows use their requested data rate q_{v_1}, q_{v_2} , and the web flow w_1 uses the remaining capacity q_r . Figure 2(b) illustrates the case in which $q_{v_1} + q_{v_2} > q_{S_v}$. The video flows share the allocated video capacity q_{S_v} according to their requested video rates.

4.3. Weighted bandwidth allocation for all video flows (WBV)

Another approach utilizes the service ratio between the two classes. The idea is that videos, in general, require a multiple of the bandwidth needed for web browsing. Thus, there is a weight parameter $w_V > 1$, such that each video flow is assigned a w_V times larger bandwidth than a web browsing flow. In contrast to FBV, there is no fixed separation between the resources of the two classes. Instead, the resource distribution between the web class and the video class scales according to the number of active flows.

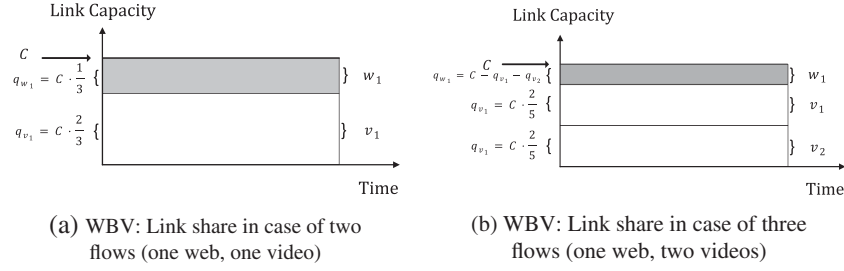


Figure 3. Behaviour of WBV for two different cases in terms of number of video flows. Weight $w_V = 2$ for both figures

Algorithm: All active flows are split into two sets, one set S_v containing all video flows and one set S_w containing all web browsing flows. Each active flow contributes to the calculation either weighted (video) or normal (web). Thus, a video flow is assigned a bandwidth of $q_v = C \cdot \frac{w_V}{w_V \cdot |S_v| + |S_w|}$, whereas a web browsing flow will be allocated $q_w = C \cdot \frac{1}{w_V \cdot |S_v| + |S_w|}$. This in turn results in a total link share with $q_{S_v} = |S_v| \cdot q_v$ and $q_{S_w} = |S_w| \cdot q_w$ as $C = q_{S_v} + q_{S_w}$.

Figure 3(a) shows the behaviour of the WBV algorithm with weight $w_V = 2$. As there are two flows sharing the link, the sharing ratio between v_1 and w_1 is 2 : 1. Figure 3(b) illustrates the case in which an additional video flow is on the link. Here, each video flow acquires 2/5 of the link capacity C , whereas w_1 acquires 1/5 of the link capacity C .

4.4. Dynamic bandwidth allocation for each video flow (DBF)

In contrast to the aforementioned described resource allocations, which are based on flow types or flow characteristics, also dynamic bandwidth allocation based on current application information is possible. In particular, the buffer fill of video flows is used in this work. Thereby, the short-term flexibility introduced by the buffer can be utilized, and struggling flows can be helped to avoid stalling, which is the worst QoE degradation for video flows. The strategy relies on an equal sharing of the capacity among all flows. Video flows that have a low buffer level (below a threshold t_l) will be added to a special set S_p of flows, which can utilize a reserved bandwidth fraction q_p . q_p is a parameter, which has to be configured by the network operator in order to adjust the extent of prioritization. This prioritization shall help to fill the buffer. If the buffer usage exceeds a threshold t_h , the video flow is removed from S_p . Note that multiple video flows may be prioritized simultaneously.

Algorithm: All active flows are put into a set S_u in the beginning. Additionally, an empty set S_p is defined for prioritized flows, and a bandwidth q_p is reserved for that set. For each video flow, the playout buffer is monitored. If the playback buffer falls below the lower threshold t_l , the flow is moved from S_u to S_p . If the playback buffer of a flow in S_p is above the high fill threshold t_h , it is moved back to S_u . All prioritized video flows share q_p equally, that is, $q_{v_p} = \frac{q_p}{|S_p|}$. All other flows (i.e. web browsing flows and non-prioritized video flows) equally share the remaining bandwidth. This means that the bandwidth allocated to non-prioritized video flows is $q_{v_u} = \frac{C - q_p}{|S_u|}$, or $q_{v_u} = \frac{C}{|S_u|}$ if S_p is empty. Web browsing flows are never prioritized and thus always receive $q_w = \frac{C - q_p}{|S_u|}$, or $q_w = \frac{C}{|S_u|}$ if S_p is empty.

Figure 4(a) illustrates the case in which the playback buffer of all videos is greater than t_l . All flows share the available capacity C equally. Figure 4(b) shows the case in which the video flow v_2 is added to the set S_p . As in this case, only the buffer of v_2 is below the threshold t_l , and v_2 receives all the allocated capacity q_p . All other flows, that is, w_1 and v_1 , share the remaining capacity $C - q_p$ equally.

5. EVALUATION METHODOLOGIES AND SCENARIO

In Section 5.1, we introduce the analytic approach that can be used to evaluate the performance of FBV, FBF and WBV. In order to analyse DBF, we provide a simulation that is described in Section 5.2.




Figure 4. Behaviour of DBF for two different cases in terms of current state of the video applications

In Section 5.3, the investigated scenario is outlined; details about the video encoding, average bit rates and so on are given.

5.1. Analytic approach

We model the access link as a birth–death process with a single server and model its utilization variations as a stationary process of singular and independent arrivals of traffic. This means that flows enter the system, they are served with a certain download bandwidth and they leave the system after the download is finished. For mathematical tractability, we consider Poisson arrivals with rate λ and a negative exponentially distributed job size (i.e. video or web page size) with mean s , which gives the service rate $\mu = \frac{C}{s}$. Thus, the model for normal TCP behaviour without any resource allocation policy, which shares the capacity equally among all flows, is described by a M/M/1 queue with processor sharing policy (M/M/1-PS). To account for different application usage, we use different arrival rates λ_v, λ_w and mean job sizes s_v, s_w for video streaming and web browsing.

This two class M/M/1-PS model is a special case of the more general discriminatory processor sharing (DPS) model described in [17]. Following the work of [16], it is possible to find the conditional average response times $V(\tau)$ for each class depending on the job size τ . Thus, for any given web page size or video size, the average download time can be given. As $V_v(\tau)$ consists of the video length (depending on video size τ and video bit rate) and the stalling time, the average stalling time should be definable. However, as we are not interested in a single user request, but in the situation in the whole network, we analyse the unconditional average response times V , for which we have

$$\begin{aligned} V_v &= \frac{1}{\mu_v(1 - \rho)} \\ V_w &= \frac{1}{\mu_w(1 - \rho)} \end{aligned} \quad (1)$$

where $\rho = \rho_v + \rho_w = \frac{\lambda_v}{\mu_v} + \frac{\lambda_w}{\mu_w}$. Note that V_v does not allow for a computation of average stalling times, as video bit rate distributions have to be taken into account. Thus, here (and for all resource allocation strategies), we follow a simulation approach as described in Section 5.2 to obtain more detailed results.

The WBV strategy corresponds to the previously described more general DPS model, in which each class is assigned a weight that controls the division of processor capacity. Thus, using the weight w for allocating w times larger bandwidth to videos than to web pages, we can use the generalized results from [16]:

$$\begin{aligned} V_v &= \frac{1}{\mu_v(1 - \rho)} \left(1 + \frac{\mu_v \rho_w (1 - w)}{w \mu_v (1 - \rho_v) + \mu_w (1 - \rho_w)} \right) \\ V_w &= \frac{1}{\mu_w(1 - \rho)} \left(1 + \frac{\mu_w \rho_v (w - 1)}{w \mu_v (1 - \rho_v) + \mu_w (1 - \rho_w)} \right) \end{aligned} \quad (2)$$

The FBV and FBF are best described by the generalized processor sharing model in which each request or class of requests receives an arbitrary service rate. Thus, FBV can be described by a model with two classes, one for video and one for web pages. From the work of [38], stationary joint distributions can be obtained, but no closed form for the conditional response times is given. However, numerical results can be obtained by solving the steady-state equations of the FBV model. FBF can be modelled by a single class for each video and one class for all web pages, which enlarges and complicates the FBV model such that we did not conduct an analytical performance evaluation. DBF is a dynamic strategy prioritizing flows depending on current playback buffers, so no analytical model can be provided. Therefore, our analysis of FBF and DBF will be based on results of the simulation runs only.

5.2. Simulation

To confirm the analytic findings on the average download times, to obtain more detailed results, and to evaluate the strategies for which no analytic approach is present, we follow a simulative approach based on the analytic model. The streaming of videos and the download of web pages over a single link is simulated with a Java discrete-event simulation. Again, a single link has a fixed capacity C , which is shared among all application flows of the users in the system. The streaming is implemented such that different flow control (i.e. bandwidth allocation) strategies can be applied. Each run simulates the link utilization slightly more than one day. The arrivals of both video and web page requests follow Poisson processes with rates λ_v and λ_w , respectively. The size of videos s_v and web pages s_w is exponentially distributed. With video streaming, the stalling of each video depends on the download bandwidth, video bit rate and buffering strategy. The average video bit rate \bar{b} is exponentially distributed. For the simulation of variable bit rate videos, the video is split in chunks of 2 s playtime. The bit rate of each chunk is different, but constant bit rate was assumed within a chunk. For the varying chunk bit rate, we use a simple auto-regressive model, which can accurately model videos without scene changes [39]. Note that, in practice, higher variations may occur if the video content is composed of several scenes, which can negatively affect the presented results. For each video, a playout buffer is maintained, and a simple buffering strategy is adopted: stalling of a video begins when the playout buffer is empty, and playback resumes after 5 s of playtime are in the buffer. Thus, not only the resulting video and web page download times but also stalling times can be obtained. In the performance evaluation, the buffering ratio, that is, the ratio of stalling time and playback time, will be used. An initial delay of 2 s was chosen as it provides some time to pre-buffer video data without significantly deteriorating the perceived quality [40].

5.3. Evaluation scenario

The aim is to investigate a recently new emerging scenario with progressive video streaming. We use the observations of [41] for our evaluations and additionally investigated the sizes of 50 randomly downloaded videos and web pages, respectively. In detail, we focus on 360p flash videos, which is the current default quality for YouTube videos on smartphones. According to [41], the average encoding bit rate of 360p YouTube videos is $\bar{b} = 0.5$ Mbps. Furthermore, we measured a mean video clip length of 110 s. Therefore, we define $s_v = \bar{b} \cdot \text{video clip length} = 6.87$ MB. To align the simulation and the analytical approach, we model the video size as a random variable following an exponential distribution with the reciprocal of s_v as rate parameter. Each video is divided into chunks with random sizes, which again follow an exponential distribution with the reciprocal of $2s \cdot \bar{b}$ as rate parameter. Only the last chunk is cropped to reach the desired video size. The size of a web page is exponentially distributed with a mean of $s_w = 1.3$ MB. For the performance evaluation, we simulate one link with capacities of 5, 10 and 20 Mbps and a maximum of 50 flows.

6. PERFORMANCE EVALUATION AND INSIGHTS

This section presents the results of the performance evaluation of all algorithms of this paper. In particular, we start with investigating for which scenarios an active resource scheduling for a link is needed,

that is, for which level of network use a resource scheduling may improve the quality of the applications. Based on this observation, the algorithms are evaluated analytically and with simulation for different traffic mixes, that is, for different sharing ratios between web and video flows for a particular network use. Finally, we provide an overview of all algorithms and differentiate their trade-offs between the quality of web and video traffic. Note that the following plots contain confidence intervals for 95% confidence level.

6.1. Situation in current networks without resource management

To motivate why an application-aware resource management is needed, the performance of all network flows in terms of QoS parameters is presented and then linked to the actually perceived quality of each application. The results allow us to state that pure network parameters may not show any performance issues, whereas looking at the application parameters reveals dramatic performance degradation.

6.1.1. Best-effort networks

As described in Section 3, it is assumed that with TCP, all flows share the available link capacity equally. This means that every flow obtains the same bandwidth share. Figure 5 shows the normalized download time of 1 MB for different capacities and network loads. Both the results of the analytical model and the simulation of Section 5 are presented. First, it can be seen that the analytical model closely resembles the simulation as the analytical model results fall within the 95% confidence intervals of the simulation. Second, for a given network load, a higher link capacity leads to lower average download times. Third, confirming the analytical model in equation (1), when the network load increases towards 1, the time to download data grows hyperbolically. This leads to unacceptable download times that are commonly avoided in most networks by over-provisioning.

6.1.2. Application performance

Although the general results of Section 6.1.1 hold, different applications have different requirements to the network. Thus, also the changing of network parameters has a different impact on the quality of the various applications. This holds especially for the two considered applications in this work, that is, video streaming and web browsing. Therefore, we analyse application inherent parameters that have a high correlation with the user-perceived quality of the application as described in Section 3.1. These key performance indicators are download time of a web page for web browsing and buffering ratio for video streaming.




Figure 5. Normalized download time of 1 MB in best-effort networks

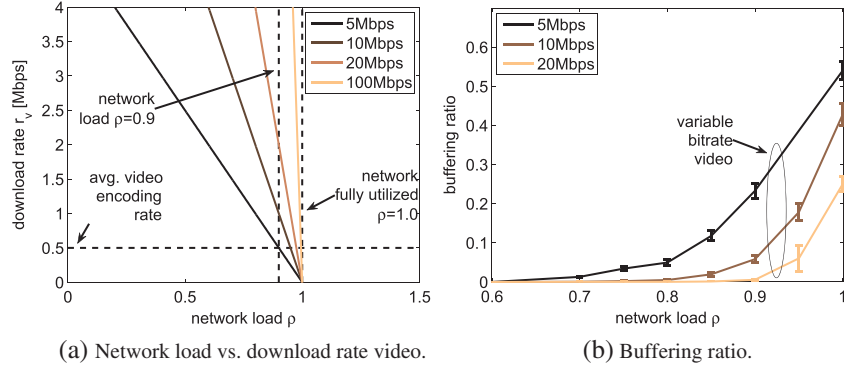


Figure 6. Video application performance in best-effort networks

Figure 6(a) compares the download rate r at different link capacities and loads for videos with an average encoding rate of 0.5 Mbps. It shows that the download rates decrease linearly when network load increases. The download rate has to be at least as high as the average video rate to guarantee a video playback without interruptions. It can be seen that a 5 Mbps link can provide such download rates up to a load of 0.9. Higher link capacities can even guarantee smooth playback for higher loads. However, this is only true for constant bit rate (CBR) videos where the download rate does not fluctuate. Currently deployed streaming applications apply pre-buffering of video content and use video encoding with VBR, which may result in different video performance. For example, with VBR, the network requirements of the video can vary over time, but the use of the buffer can overcome certain load peaks. Therefore, this behaviour was emulated in the previously presented simulation as described in Section 5.2. Figure 6(b) shows the simulation results for the average buffering ratio of VBR videos for different capacities and network loads. It can be seen that, in contrast to CBR, stalling of videos with slightly varying bit rate can occur already with loads around 0.7 and above. Consequently, the benefits of resource allocation strategies in high load scenarios will be investigated in the following.

6.2. Performance evaluation of resource management algorithms

This section evaluates the resource management algorithms, which were presented in Section 4. In this section, if not stated otherwise, results for a link capacity of 10 Mbps and a network load of 0.9 are shown. Note that this is the load of a best-effort system (M/M/1-PS) with given arrival rates λ_v, λ_w , job sizes s_v, s_w and capacity C (cf. Section 5), which is also used for the systems that apply the different resource allocation algorithms. Nevertheless, this general approach can also evaluate situations with different network parameters.

This analysis investigates the key QoE influence factors of the two considered applications, web browsing and video streaming. As the performance of a given resource management algorithm depends not only on application characteristics but also on their popularity in the network, the traffic mix ($\frac{\lambda_v}{\lambda_w}$) is introduced. For example, a traffic mix of $\frac{\lambda_v}{\lambda_w} = 2$ means that on average, there are two times more video downloads than web page downloads over the link.

6.2.1. FBV

The FBV strategy reserves a certain share of the capacity to video flows. The challenge, however, is the selection of the right prioritization settings. If too much bandwidth is allocated for video, the download time of web pages is considerably affected. If the reserved rate is too small, the video may stall and the prioritization is useless. Therefore, several fixed allocations ranging from $q_{S_v} = 0.3C$ up to $q_{S_v} = 0.9C$ are simulated. Figure 7(a) depicts the analytical and simulative results of the average download time of a web page for a fixed network load $\rho = 0.9$ and varying traffic mixes. It can be seen that the download times in a best-effort network are 9 s on average for each traffic mix. On the

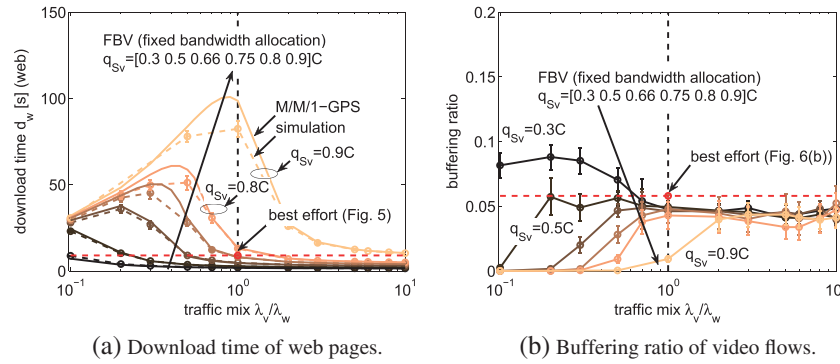


Figure 7. FBV: Fixed bandwidth reservation q_{sv} for all video traffic at a load $\rho = 0.9$

left side of the dashed vertical line, more web page downloads are on the link. The plot depicts that an increased bandwidth allocation generally increases the web page download times as videos are in the system, which leave little resources for the web page downloads. When the traffic mix increases towards $10^0 = 1$ and beyond, this effect is mitigated because less web pages are in the system. Thus, they are not penalized to such a large extent as the bandwidth share becomes more aligned to the traffic share. In the rightmost part of the figure, only few web flows are in the system. They can then use the remaining (i.e. not reserved for videos) capacity almost exclusively, which stems from but also leads to short download times. For the analytical M/M/1-GPS model, the mean number of users in the system is obtained from the state probabilities. Little's theorem is used to compute the mean delay in the system, which is equal to the mean service time in a processor sharing system. As can be seen in the plot, the simulation runs are approximated very well by the analytical model. Thus, also the analytical model can be used for the performance evaluation.

Figure 7(b) shows the exactly same situation from the perspective of video flows and depicts the average buffering ratio for different traffic mixes and different bandwidth allocations. Again the best-effort performance (cf. Figure 6(b)) is shown as a base value. In the rightmost part of the figure, it can be seen that the buffering ratio converges towards the best-effort performance for any prioritization. This can be explained by the fact that web flows are seldom in the system and for short times only. This means that there are many videos that compete most of the time only among themselves, which leads to a best-effort behaviour. In the left part of the plot, video flows are more seldom in the system but have no problems because of the bandwidth allocation. In general, a higher bandwidth allocation leads to smaller buffering ratio with traffic mixes below 1. Only in the case of $q_{sv} = 0.3C$ is the allocation not sufficient as the buffering ratio is higher than the best-effort base line. This again shows the need for careful selection of the right allocation parameter. The trade-off between the optimization of the two applications will be further discussed in Section 6.3.

6.2.2. FBF

The difference of FBF compared with FBV is that it allocates bandwidth to individual video flows. As described in Section 4, two allocations are investigated. First, each video is assigned its average bit rate, and second, each video is assigned its average bit rate plus standard deviation to account for bit rate variations (VBR content). Figure 8(a) shows the average download times of web pages for both allocations. For comparison, selected FBV performances ($q_{sv} = [0.3, 0.6, 0.9]C$) are depicted in grey. Obviously, it can be observed that the FBF download times increase if the traffic mix increases. This is because more and more video flows are in the system, which request bandwidth; hence, the reserved bandwidth for video, which is unavailable for web page downloads, adds up.

Figure 8(b) shows the buffering ratio of the video flows in the same scenario. It can be seen that the average video bit rate allocation is not sufficient as it cannot avoid stalling of VBR videos. However, the allocation of average video rate plus standard deviation performs quite well for traffic mixes below 1 where it almost completely avoids stalling. When more video flows are in the system, the reservation

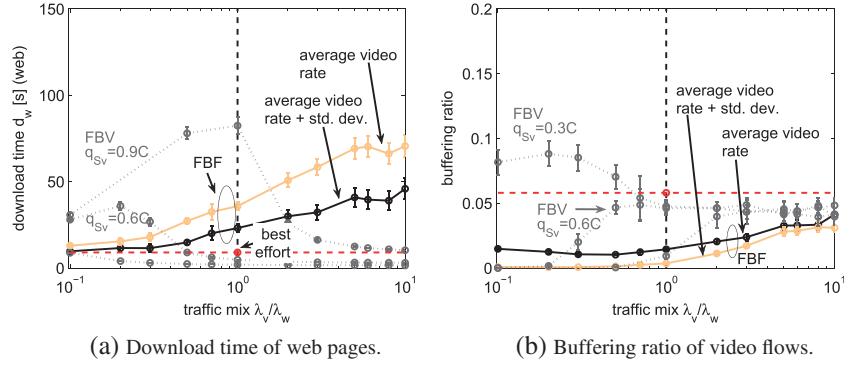


Figure 8. FBF: fixed bandwidth reservation for each flow with either average bit rate or average bit rate plus standard deviation at link load $\rho = 0.9$

of the videos exceeds the available capacity. Thus, each video receives a share relative to its request, which leads to increasing buffering ratios and consequently bad performances of the video streams. With increasing traffic mix, it performs better than all FBV algorithms, which use an equal share policy if too many videos are in the system, as the relative sharing of bandwidth is better aligned to the demands of the video flows.

6.2.3. WBV

The WBV strategy assigns a weight w_v to each video flow such that it receives a w_v times larger bandwidth than a web flow. Figure 9(a) shows the results of the analytic DPS model (cf. equation (2) in Section 5) and of the simulation runs. It can be seen that the analytic model and the simulation results are well aligned, especially for smaller values of w_v . Moreover, a larger w_v increases the download time, as web flows receive less bandwidth compared with video flows for any traffic mix. Only for large weights and large traffic mixes does the bandwidth given to web flows become almost arbitrarily small, which leads to a fast increase of download times.

The buffering ratio in the same scenario can be seen in Figure 9(b). The video flows benefit from WBV with $w_v > 1$ in case of traffic mix below 1. If more video flows are in the system, the effect reduces and the buffering ratio again converges towards that of best-effort networks (equal sharing). It can be seen that WBV is suitable to align bandwidth allocation to the flow sizes by choosing an appropriate weight w_v for the given traffic mix. As the resulting resource allocation also accounts for the number of flows, it can significantly reduce buffering ratios for video flows while only slightly increasing the download time of web pages. Again, trade-offs between the two applications have to be taken into account and are further investigated in Section 6.3.

6.2.4. DBF

The results presented so far showed the performance of resource allocation algorithms in which a parameter was set statically for each flow. This means the bandwidth allocation changed only depending on the current number of flows and the respective parameters. In contrast, DBF is a dynamic resource allocation strategy, which prioritizes a video flow if its buffer fill is below a threshold t_l and de-prioritizes it once its buffer level has passed a threshold t_u . For the performance evaluation of DBF, the prioritization bandwidth q_p , which can be utilized by suffering video flows, is set to $q_p = 0.9 \cdot C$, and the thresholds are set to $t_l = 10$ s and $t_u = 20$ s. Figure 10(a) shows the performance of DBF (dashed) and compares it with WBV (solid), which showed a decent performance with respect to both traffic classes. It can be seen that the download time of web pages is highest for an equal traffic mix. The more the traffic mix becomes imbalanced (either more videos or more web pages are in the system), the download time improves and comes close to the best-effort situation.

The video flows, on the other hand, face a significant performance improvement, which is depicted in Figure 10(b). If there are many video flows in the system, the buffering ratio is worst but still converges

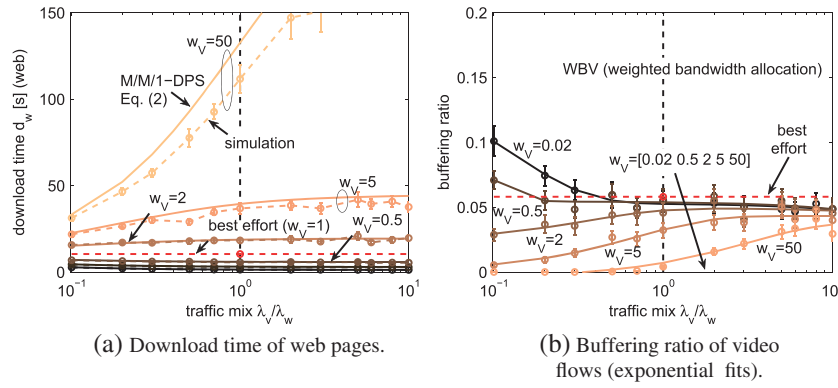


Figure 9. WBV: weighted bandwidth reservation for all video traffic with weight w_V at link load $\rho = 0.9$

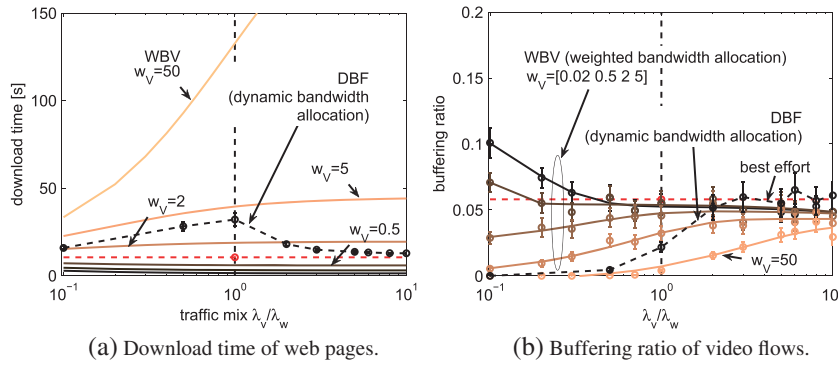


Figure 10. DBF: dynamic bandwidth allocation for each video flow based on buffer fill at link load $\rho = 0.9$

to the best-effort situation. However, if there are less videos in the system, the buffering ratio is much smaller than without resource allocation strategy. It can be seen that DBF supports video flows even better than WBV with a high $w_V > 5$ but at the same time does not penalize web flows too much.

To sum up, DBF has been shown to effectively reduce stalling for traffic mixes $\frac{\lambda_v}{\lambda_w} < 2$. But although DBF is a video prioritization algorithm, it only slightly increases the download time of web pages. Thus, DBF ranges among the best investigated algorithms especially taking balanced and fair behaviour into account. However, it has to be noted that the influence of the parameters q_p , t_l and t_u was not investigated in this work and is thus still open. Moreover, the gain of dynamic allocation based on application state comes with the cost for monitoring and signalling application information to the resource management entity. The trade-offs for this additional cost will also be investigated in future work.

6.3. Comparison of the different strategies

After presenting the performance of each individual algorithm, we conduct a Pareto analysis to compare how the different strategies influence the QoE in our scenario. Therefore, the two QoE performance indicators, that is, download time for web flows and buffering ratio for video streams, are used as axes for the Pareto plot in Figure 11. The best-effort performance (mean download time 9.04 s, mean buffering ratio 0.058) is depicted by a red cross and serves as a reference. Similarly, the Pareto-optimal performances of the algorithms for a given traffic mix are marked in the plot. We distinguish five different traffic mixes by coloured lines ranging from 0.1 (black), that is, ten times more web flows

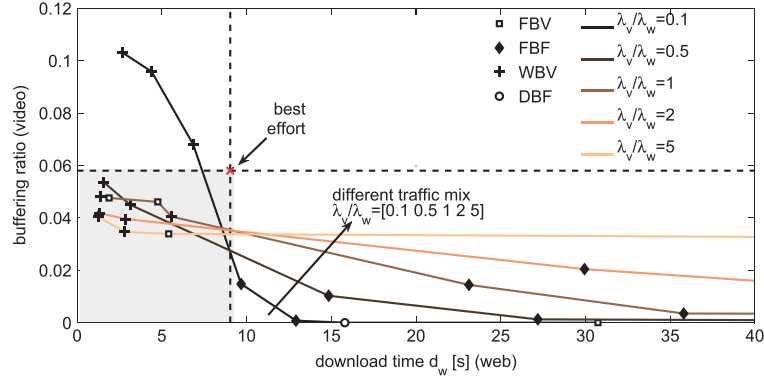


Figure 11. Overview of Pareto-optimal performances for different traffic mixes

than video flows, over 1 (light brown), that is, the same amount of web and video flows, to 5 (yellow), that is, five times more video flows than web flows. Note that only the marked points are Pareto-optimal algorithm performances; the lines connecting the points, which belong to the same traffic mix, however, are only for visualization purposes and do not indicate the location of other Pareto optimal. The marker shape of the Pareto-optimal points represents the type of algorithm, which achieved the respective results.

The Pareto-optimal performances, that is, performances, in which it is impossible to improve in one dimension without deteriorating the other dimension, are either below or left of the best-effort performance. This means the usage of any of the resource allocation algorithms does improve the performance at least in one dimension. It can be seen that especially WBV algorithms can be used to optimize the web page download time, whereas FBF and DBF achieve minimal buffering ratios. The most interesting part is the bottom left grey shaded box, which indicates an improvement for both web and video flows. The marker shapes indicate that WBV and FBF algorithms can achieve such performances when they use the right parameter for the given traffic mix. Note that the specific algorithm parameter for a Pareto optimum is not given in this plot but can be obtained from the respective plots in the previous sections.

7. CONCLUSION

This work presented new methods that allow analysing application-aware resource management analytically and with simulation. Application-aware resource management aims at improving not only network's resource efficiency but also users' satisfaction, that is, the application QoE of users. To count for varying traffic situations, a joint optimization of two application types, namely, web browsing and video streaming, is considered, which has not yet been investigated by the research community. With the proposed methods, network operators can quantify performance gain trade-offs between investing in pure capacity increase (over-provisioning) and the realization of application-aware resource management approaches.

This work proposed a model for the application-aware resource management problem and discussed four algorithms for resource management. The algorithms vary in complexity and implementation levels in order to cover the design space of application-aware resource management. In order to evaluate the performance of the algorithms, analytical models based on processor sharing queues (M/M/1-PS) were presented if possible. Moreover, a discrete-event simulation was implemented in Java to assess all algorithms.

Furthermore, a first study of application-aware resource management approaches of this paper provides configuration guidelines for the selection and configuration of the considered algorithms. The results of the study show that it is possible to improve the quality of all applications compared with best-effort transmission. This means that for specific traffic mixes, it is possible to obtain a significant

gain in network resource use compared with simple best-effort transmission while still having a low implementation complexity. In case of more complex, that is, more dynamic algorithms, the application quality of videos was significantly improved. Based on all analysis, operators have the possibility to set up a resource management according to their objectives and traffic situations of their networks. Such an analysis in terms of quality of experience and varying traffic conditions has not yet been considered in the literature.

REFERENCES

1. Ferguson AD, Guha A, Liang C, Fonseca R, Krishnamurthi S. Participatory networking. In *ACM SIGCOMM 2013 Conference—SIGCOMM '13*. ACM Press: New York, New York, USA, 2013, pp. 327.
2. Georgopoulos P, Elkhatib Y, Broadbent M, Mu M, Race N. Towards network-wide QoE fairness using OpenFlow-assisted adaptive video streaming. In *ACM SIGCOMM 2013 Workshop on Future Human-Centric Multimedia Networking—FHMN '13*. ACM Press: New York, New York, USA, 2013, pp. 15.
3. Huang T, Johari R, McKeown N. Downton abbey without the hiccups. In *ACM SIGCOMM 2013 Workshop on Future Human-Centric Multimedia Networking—FHMN '13*. ACM Press: New York, New York, USA, 2013, pp. 9.
4. Jain S, Kumar A, Mandal S, Ong J, Poutievski L, Singh A, Venkata S, Wanderer J, Zhou J, Zhu M, Zolla J, Hözl U, Stuart S, Vahdat A. B4: Experience with a globally-deployed software defined WAN. In *ACM SIGCOMM 2013 Conference—SIGCOMM '13*. ACM: Hong Kong, China, 2013, pp. 3–14.
5. Jarschel M, Wamser F, Höhn T, Zinner T, Tran-Gia P. SDN-based application-aware networking on the example of YouTube video streaming. In *2nd European Workshop on Software Defined Networks (EWSN 2013)*, Berlin, Germany, 2013; 87–92.
6. Qazi ZA, Lee J, Jin T, Bellala G, Arndt M, Noubir G. Application-awareness in SDN. In *ACM SIGCOMM 2013 Conference—SIGCOMM '13*. ACM: New York, NY, USA, 2013; 487–488.
7. Thakolsri S, Khan S, Steinbach E, Kellerer W. QoE-driven cross-layer optimization for high speed downlink packet access. *Journal of Communications, Special Issue on Multimedia Communications, Networking and Applications* 2009; **4**(9): 669–680.
8. Tian G, Liu Y. Towards agile and smooth video adaptation in dynamic HTTP streaming. In *Proceedings of the 8th International Conference on Emerging Networking Experiments and Technologies—CoNEXT '12*. ACM Press: New York, New York, USA, 2012, pp. 109.
9. Wamser F, Hock D, Seufert M, Staehle B, Pries R, Tran-Gia P. Using buffered playtime for QoE-oriented resource management of YouTube video streaming. *Transactions on Emerging Telecommunications Technologies* 2013; **24**(3): 288–302.
10. Khan S, Peng Y, Steinbach E, Sgroi M, Kellerer W. Application-driven cross-layer optimization for video streaming over wireless networks. *IEEE Communications Magazine* 2006; **44**(1): 122–130.
11. Wamser F, Zinner T, Iffländer L, Tran-Gia P. Demonstrating the prospects of dynamic application-aware networking in a home environment. In *ACM SIGCOMM 2014, Demonstrations*, Chicago, IL, USA, August 2014; 149–150.
12. Staehle B, Hirth M, Pries R, Wamser F, Staehle D. YoMo: A YouTube application comfort monitoring tool. In *New Dimensions in the Assessment and Support of Quality of Experience for Multimedia Applications*, Tampere, Finland, 2010; 1–3.
13. Wamser F, Zinner T, Zhu J, Tran-Gia P. Dynamic bandwidth allocation for multiple network connections: Improving user QoE and network usage of YouTube in mobile broadband. In *ACM SIGCOMM Capacity Sharing Workshop (CSWS 2014)*, Chicago, IL, USA, August 2014; 57–62.
14. Hock D, Wamser F, Seufert M, Pries R, Tran-Gia P. OC²E²AN: Optimized control center for experience enhancements in access networks. *PIK—Praxis der Informationsverarbeitung und Kommunikation* 2013; **36**: 40.
15. Superiori L, Wrulich L, 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, 2009. IWCLD'09*. IEEE: Palma de Mallorca, Spain, 2009; 1–5.
16. Fayolle G, Mitran I, Iasnogorodski R. Sharing a processor among many job classes. *Journal of the ACM (JACM)* 1980; **27**(3): 519–532.
17. Kleinrock L. Time-shared systems: A theoretical treatment. *Journal of the ACM (JACM)* 1967; **14**(2): 242–261.
18. Garcia-Dorado JL, Finamore A, Mellia M, Meo M, Munafo M. Characterization of ISP traffic: trends, user habits, and access technology impact. *IEEE Transactions on Network and Service Management* 2012; **9**(2): 142–155.
19. Maier G, Feldmann A, Paxson V, Allman M. On dominant characteristics of residential broadband internet traffic. In *ACM SIGCOMM 2009 Conference on Internet Measurement—IMC '09*. ACM Press: New York, New York, USA, 2009, pp. 90.
20. Paul S, Jain R. OpenADN: Mobile apps on global clouds using OpenFlow and software defined networking. In *Globecom Workshops (GC Wkshps), 2012 IEEE*. IEEE: Anaheim, CA, USA, 2012; 719–723.
21. Nordström E, Shue D, Gopalan P, Kiefer R, Arye M, Ko S, Rexford J, Freedman MJ. Serval: An end-host stack for service-centric networking. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI 12)*, San Jose, CA, USA, 2012; 85–98.
22. Martin J, Feamster N. User-driven dynamic traffic prioritization for home networks. In *ACM SIGCOMM 2012 Workshop on Measurements up the Stack (W-MUST'12)*. ACM: Helsinki, Finland, 2012; 19–24.
23. Katchabaw MJ, Lutfiyya HL, Bauer MA. Usage based service differentiation for end-to-end quality of service management. *Computer Communications* 2005; **28**(18): 2146–2159.

24. 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, May 2014; 1–6.
25. Zinner T, Hößfeld T, Fiedler M, Liers F, Volkert T, Khondoker R, Schatz R. Requirement driven prospects for realizing user-centric network orchestration. *Multimedia Tools and Applications* 2014; **74**(2): 413–437.
26. Huang C, Juan HH, Lin MS, Chang CJ. Radio resource management of heterogeneous services in mobile WiMAX systems [Radio Resource Management and Protocol Engineering for IEEE 802.16]. *IEEE Wireless Communications* 2007; **14**(1): 20–26.
27. Saul A. Simple optimization algorithm for MOS-based resource assignment. In *VTC Spring 2008—IEEE Vehicular Technology Conference*. IEEE: Singapore, May 2008; 1766–1770.
28. Reis AB, Chakareski J, Kassler A, Sargento S. Quality of experience optimized scheduling in multi-service wireless mesh networks. In *IEEE Conference on Image Processing (ICIP)*. IEEE: Hong Kong, China, 2010; 3233–3236.
29. Tang P, Wang P, Wang N, Nguyen Ngoc V. QoE-based resource allocation algorithm for multi-applications in downlink LTE systems. In *Proceedings of the 2014 International Conference on Computer, Communications and Information Technology* Atlantis Press: Paris, France, 2014, pp. 88–93.
30. IETF Working Group on Application-Layer Traffic Optimization (ALTO).
31. Alimi R, Yang Y, Penno R. *Application-Layer Traffic Optimization (ALTO) protocol*, September 2014.
32. Heller B, Seetharaman S, Mahadevan P, Yiakoumis Y, Sharma P, Banerjee S, McKeown N. Elastictree: Saving energy in data center networks. In *USENIX Conference on Networked Systems Design and Implementation (NSDI'10)*, Vol. 10, San Jose, CA, USA, 2010, pp. 249–264.
33. Jiang JW, Lan T, Ha S, Chen M, Chiang M. Joint VM placement and routing for data center traffic engineering. In *IEEE INFOCOM 2012*. IEEE: Orlando, FL, USA, 2012; 2876–2880.
34. Xie D, Ding N, Hu YC, Kompella R. The only constant is change: Incorporating time-varying network reservations in data centers. *ACM SIGCOMM Computer Communication Review* 2012; **42**(4): 199–210.
35. Egger S, Reichl P, Hößfeld T, Schatz R. Time is Bandwidth? Narrowing the gap between subjective time perception and quality of experience. In *2012 IEEE International Conference on Communications (ICC 2012)*, Ottawa, Canada, June 2012; 1325–1330.
36. Hößfeld T, Schatz R, Seufert M, Hirth M, Zinner T, Tran-Gia P. Quantification of YouTube QoE via crowdsourcing. In *IEEE International Workshop on Multimedia Quality of Experience—Modeling, Evaluation, and Directions (MQoE 2011)*, Dana Point, CA, USA, December 2011; 494–499.
37. Dobrian F, Awan A, Joseph D, Ganjam A, Zhan J, Sekar V, Stoica I, Zhang H. Understanding the impact of video quality on user engagement. In *ACM SIGCOMM Conference*. ACM: Toronto, ON, Canada, 2011; 362–373.
38. Cohen JW. The multiple phase service network with generalized processor sharing. *Acta Informatica* 1979; **12**(3): 245–284.
39. Tanwir S, Perros H. A survey of VBR video traffic models. *IEEE Communications Surveys & Tutorials* 2013; **15**(4): 1778–1802.
40. Hößfeld T, Egger S, Schatz R, Fiedler M, Masuch K, Lorentzen C. Initial delay vs. interruptions: Between the devil and the deep blue sea. In *QoMEX 2012*, Yarra Valley, Australia, July 2012; 1–6.
41. Finamore A, Mellia M, Munafò MM, Torres R, Rao SG. YouTube everywhere: Impact of device and infrastructure synergies on user experience. In *ACM SIGCOMM 2012 Conference on Internet Measurement Conference—IMC '11*. ACM Press: New York, New York, USA, 2011, pp. 345.

AUTHORS' BIOGRAPHIES

Florian Wamser is a post-doctoral fellow at the Chair of Communication Networks at the University of Würzburg of Prof. Dr-Ing. P. Tran-Gia. He heads the Working Group on Network Resource Management and Quality of Experience. After leaving school, he founded a company for the design and development of educational computer networks, which he later gave up to study computer science in Würzburg and Helsinki, Finland. His PhD thesis is entitled “Performance Assessment of Resource Management Strategies in Cellular and Wireless Access Networks” and deals with network resource management that takes into account the specific requirements of applications on the network. His current research focus is on analytical and simulative performance evaluation and optimization of wireless broadband access networks and related fields as well as cellular communication. He is the main developer of YoMo, a program that is able to measure the satisfaction of a user with progressive streaming video (such as YouTube) and designed a series of demonstrations to show the benefits of a joint network and application control in a home or wireless access networks.

Andreas Blenk studied Computer Science at the University of Würzburg, Germany, where he received his diploma degree in 2012. After this, he joined the Chair of Communication Networks at the Technische Universität München in June 2012, where he is working as a research and teaching associate. As a member of the software-defined networking and network virtualization research group, he is currently pursuing his PhD. His research is focused on service-aware network virtualization, virtualizing software-defined networks, as well as resource management and embedding algorithms for virtualized networks.

Michael Seufert studied Computer Science, Mathematics and Education at the University of Würzburg, Germany. In 2011, he received his diploma degree in Computer Science and additionally passed the first state examinations for teaching Mathematics and Computer Science in secondary schools. From 2012–2013, he has been with FTW Telecommunication Research Center in Vienna, Austria, working in the area of user-centered interaction and communication economics. Currently, he is a researcher at the Chair of Communication Networks of the University of Würzburg and pursuing his PhD. His research mainly focuses on QoE of Internet applications, social networks, performance modelling and analysis and traffic management solutions.

Thomas Zinner studied Computer Science at the University of Würzburg, Germany. He finished his PhD on performance modelling of QoE-aware multipath video transmission in the future Internet in 2012. He is now heading the “Next Generation Networks” research group at the Chair of Communication Networks in Würzburg. His main research interests cover video streaming, QoE-aware networking, SDN and NFV, as well as the performance assessment of these technologies and architectures.

Wolfgang Kellerer is a full professor at the Technische Universität München (TUM) heading the Chair of Communication Networks in the Department of Electrical and Computer Engineering since 2012. Before, he has been a director and head of wireless technology and mobile network research at NTT DOCOMO’s European research laboratories, DOCOMO Euro-Labs, for more than 10 years. His research focuses on concepts for the dynamic control of networks (Software Defined Networking), network virtualization and network function virtualization and application aware traffic management. In the area of wireless networks, the emphasis is on machine-to-machine communication, device-to-device communication and wireless sensor networks with a focus on resource management towards a concept for fifth generation of mobile communications (5G). His research resulted in more than 200 publications and 29 granted patents in the area of mobile networking and service platforms. He is a member of ACM, VDE ITG and a senior member of IEEE.

Phuoc Tran-Gia is a professor and director of the Chair of Communication Networks, University of Würzburg, Germany. He is also member of the Advisory Board of Infosim (Germany) specializing in IP network management products and services. Prof. Tran-Gia is also a cofounder and board member of Weblabcenter Inc. (Dallas, Texas), which specialized in Crowdsourcing technologies. Previously, he was at Academia in Stuttgart, Siegen (Germany), as well as at Industries at Alcatel (SEL) and IBM Zurich Research Laboratory. He is active in several EU framework projects and COST actions. Prof. Tran-Gia was a coordinator of the German-wide G-Lab Project “National Platform for Future Internet Studies” aiming to foster experimentally driven research to exploit future Internet technologies. His research activities focus on performance analysis of the following major topics: Future Internet and Smartphone Applications; QoE Modeling and Resource Management; Software Defined Networking and Cloud Networks; Network Dynamics and Control; and Crowdsourcing. He has published more than 100 research papers in major conferences and journals and received the Fred W. Ellersick Prize 2013 (IEEE Communications Society).