QOE-DRIVEN RESOURCE OPTIMIZATION FOR USER GENERATED VIDEO CONTENT IN NEXT GENERATION MOBILE NETWORKS

Ali El Essaili[1], Eckehard Steinbach[1], Daniele Munaretto[2], Srisakul Thakolsri[3], Wolfgang Kellerer[3]

[1]Technische Universität München, Institute for Media Technology
[2]University of Padova
[3]DOCOMO Communications Laboratories Europe GmbH

ABSTRACT

The increasing popularity of user-generated content and the high quality upstreaming capabilities of mobile phones indicate a prevalence of video traffic in the uplink of next generation mobile networks. Need arises for optimizing the network resource allocation while preserving the user satisfaction. In this paper, we propose a service-centric approach for uplink distribution of real-time user-generated content based on the Quality of Experience (QoE) and popularity of the video content. In case of limited network resources, the proposed approach assigns more resources for popular contents while maintaining a minimum guaranteed QoE for the less popular ones. We compare our service-centric approach with a QoE-driven one that does not consider video popularity and evaluate both approaches for the uplink of an LTE system. The simulation results show that a significant gain in terms of average user satisfaction can be achieved.

Index Terms— Quality of Experience, service-centric resource allocation, LTE uplink, video popularity

1. INTRODUCTION

There is a proliferation of mobile phones equipped with digital cameras that allow the upstreaming of high quality multimedia content. Users capture real-time events and share them with other users, for instance, on video portals. The analysis of large-scale User-Generated Content (UGC) shows that the users' requests are highly skewed towards popular videos [1]. In their work on YouTube traffic characterization, the authors of [2] find that the video popularity is Zipf-like.

In this paper, we try to answer the following question: Given that not all upstreamed videos have the same popularity, can we improve the user satisfaction by optimizing the uplink resource allocation for the live captured videos? Despite the growing interest in UGC systems and services, there is no prior work that addresses our problem. Popularity has been traditionally exploited in cache management for proxy servers, whereby a proxy stores the initial frames of popular videos [3]. In [4], an analytical model for the design of coding strategies for time-shifted personalized video content is described. The authors show that a higher multicast gain can be achieved by considering the content popularity. In [5], the skewed popularity distribution in file sharing is utilized for optimal placement of the resources in structured Peer-to-Peer networks. By accelerating the search for popular contents, the average search cost for the whole system is reduced. Most relevant for our work is the popularity-aware scheduling for network coding based content distribution in ad hoc networks [6]. Network-coded blocks are assigned a popularity value based on the requests from neighboring nodes. The

transmission efficiency can be improved by assigning higher channel access priorities to popular blocks. Class-based resource allocation has been intensively studied in the literature (e.g., [7]). This is fundamentally different from our work as we study the popularity of UGC within the same video streaming service class.

Different from previous works, we propose a method for uplink distribution of live video contents by addressing the popularity of the content (i.e., number of followers), the video characteristics, and the available network resources. We introduce a service-centric concept that is based on video consumer-producer coordination through a video portal (Figure 1). In such a scenario, multiple users are simultaneously connecting to a video portal for sharing their captured live video content. The video portal ranks the video contents based on the consumers' requests and provides a central entity in the operator's network (e.g., eNodeB) with a standardized feedback about the popularity of the videos. This entity is then responsible for scheduling and resource allocation among multiple video producers. We see a big potential gain by optimizing the uplink resources based on the popularity of video contents. Knowing that not all the live content uploaded to the portal has the same number of followers, our approach allows for better user satisfaction among video consumers and provides efficient use of the wireless medium.

The paper is organized as follows. In the next section we give an overview of QoE-driven uplink resource allocation. Section 3 describes our service-centric model for the LTE uplink. In Section 4 we present our simulation results and Section 5 concludes this paper.

2. QOE-DRIVEN UPLINK RESOURCE ALLOCATION

The uplink of next generation mobile networks has to cope with community portals for the upload and upstream of live media data. As a result, new mechanisms for optimizing the video delivery are required to match the user expectations. Cross-layer approaches for resource allocation are considered to address this issue by exchanging information across the different protocol layers [8]. More specifically, application-driven cross-layer optimization (CLO) maximizes the user satisfaction by jointly optimizing the application layer, link layer and physical layer [9]. Recently, Quality of Experience (QoE) based CLO has been proposed for resource allocation in wireless networks (e.g., [10]).

In this paper, we consider a QoE-driven approach for uplink resource allocation. In particular, we are interested in characterizing the QoE in video streaming applications. The utility function for video streaming is defined in [10] by: $U = MOS(R)$, where MOS is the Mean Opinion Score (MOS) function and R is the transmission data rate. We assume a linear mapping between MOS and Peak Signal to Noise Ratio (PSNR). MOS can take on any value between
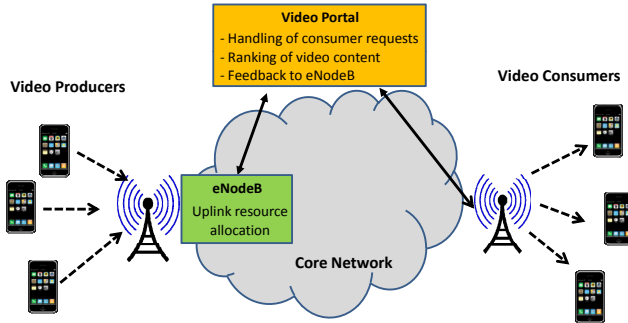
Fig. 1. Schematic depiction of the proposed service-centric approach: A video portal collects the consumers' requests and provides eNodeB with a feedback on the popularity of video contents; eNodeB allocates the uplink resources among the video producers.

Fig. 2. Mean Opinion Score (MOS) as a function of data rate.

1.0 (30 dB) and 4.5 (42 dB), which represent the worst and best QoE, respectively.

In a CLO context, varying the transmission rate at the radio link layer allows the video application to adjust its encoding parameters (e.g., rate, quantization parameter) to maximize its utility function. Different from the downlink where the video application is constrained by the applied transcoding mechanism, an arbitrary set of encoding parameters can be defined in the uplink (i.e., at the video encoder). Figure 2 shows the utility functions for 10 different video sequences encoded with the H.264 AVC video codec at QCIF resolution and a frame rate of 30 frames/sec. The application model from [11] is used to generate an arbitrary set of points for each sequence. Each video sequence exhibits a different MOS-Rate granularity. In a multi-user scenario where each user is upstreaming a different video, the transmission rates of the different users can be determined such that the overall QoE is maximized.

3. SERVICE-CENTRIC MODEL FOR LTE UPLINK

3.1. Service-centric resource allocation

We extend the QoE-driven resource allocation into a service-centric one by incorporating the feedback from a live video portal into the uplink optimization. We consider that the popularity of a video content follows a Zipf-Mandelbrot law [12]. The popularity of a content of rank k out of a population of C contents is defined by:

$$p_k = \frac{1/(k+q)^s}{H_{C,q,s}}, H_{C,q,s} = \sum_{i=1}^{C} \frac{1}{(i+q)^s} \qquad (1)$$

where q and s are the shift and shape parameters of the distribution, respectively. By setting s to 0, all contents have the same popularity. As s increases, more requests are made for popular contents. The content popularity, p_k, is a weighting factor that shapes the utility of user k according to the importance of the uploaded content. Given a population of N video upstreaming users, each uploading one content at a time, the utility-based maximization is defined by:

$$\tilde{\boldsymbol{x}}_{opt} = \arg\max_{\tilde{\boldsymbol{x}} \in \tilde{\boldsymbol{X}}} \sum_{k=1}^{N} U_k(\tilde{\boldsymbol{x}}).p_k \quad \text{where} \sum_{k=1}^{N} p_k = 1 \qquad (2)$$

$\tilde{\boldsymbol{x}}_{opt}$ and $\tilde{\boldsymbol{X}}$ are the optimal and the set of possible optimization parameters abstracted from the protocol layers, respectively [13].

Please note that the formulation in (2) maximizes the sum of all users objective functions. Alternative formulations are also possible (e.g., max-min utility [14]).

3.2. LTE uplink model

To evaluate our service-centric approach we developed an LTE uplink simulator which follows the 3GPP LTE recommendations [15]. Our resource allocation problem is to assign the physical resource blocks (PRBs) to different users such that the overall QoE is maximized. We consider a long-term radio link layer model with optimization periods in the order of seconds. This allows us to integrate our approach with any of the state-of-the art schedulers for LTE uplink. Our objective is thus to determine the amount of resources (i.e., PRBs) assigned to each user in each optimization cycle. Once the resource share for each user is determined, the deployed scheduler can assign the blocks in a contiguous manner.

To determine the achievable throughput per PRB for a given Signal-to-Noise ratio (γ) we use the LTE link layer model from [16]. The model approximates the throughput in the uplink, after link adaptation and hybrid automatic repeat request (HARQ), by an attenuation factor $\beta = 0.4$ compared to the Shannon capacity (3). As baseline uplink parameters, [16] defines a γ_{min} of -10 dB, a γ_{max} of 15 dB and a maximum throughput Thr_{max} of 2 bps/Hz.

$$Thr = \begin{cases} 0 & \text{for } \gamma < \gamma_{min} \\ \beta \log_2(1+\gamma) & \text{for } \gamma_{min} < \gamma < \gamma_{max} \\ Thr_{max} & \text{for } \gamma > \gamma_{max} \end{cases} \qquad (3)$$

We consider that no instantaneous channel quality indicator (CQI) is available but rather a long-term CQI update for each user. At each optimization cycle, we generate a new channel realization that follows the typical CQI distribution of the users in an urban macrocell [17]. For a summary of the LTE parameters, please refer to Table 1.

3.3. Cross-layer model

To abstract the application parameters (i.e., utility) and the link layer parameters (i.e., data rate) we use the long-term cross-layer abstraction from [14]. The model defines the data rate R_k for user k as a function of its resource share α_k and its maximum achievable rate $R_{max,k}$ if all the PRBs are allocated exclusively to user k, cf. (4).

Table 1. LTE Parameters

Parameter	Value
System bandwidth	5 MHz
Number of PRBs	25
Number of subcarriers	300
Bandwidth per PRB	180 KHz
Link layer model [16]	see (3)
Channel model	Urban macrocell [17]
CQI averaging cycle	1 sec
Simulation time	30 sec

Table 2. Simulation Parameters

Parameter	Value
Num of sequences	25
Num of upstream users	5...50
Num of downstream users	1000
Zipf shape parameter	1
PSNR-MOS mapping	Linear:(1,30 dB),(4.5,42 dB)
Application type	Video streaming
PSNR-Rate model	from [11]
Simulation runs	200

$$R_k = f_k(\alpha_k) = \alpha_k R_{max,k} \quad 0 \leq \alpha_k \leq 1, \forall k \qquad (4)$$

As a measure of utility, we use the MOS as defined in Section 2. Our utility function in (2), that maximizes the sum of utilities of N users given each's content popularity, can then be described by:

$$\arg\max_{(\alpha_1,...,\alpha_N)} \sum_{k=1}^{N} MOS_k(R_k).p_k \quad \text{subject to} \sum_{k=1}^{N} \alpha_k = 1 \qquad (5)$$

Each α_k value corresponds to the fraction of total PRBs assigned to user k. A greedy algorithm, similar to the work in [18], is developed to determine the value of α_k. It is initialized by assigning equal amount of resources to every user. The algorithm iteratively takes a small amount of resources from the user who is the least sensitive to the decrease in resources and assigns it to the user who gets the maximum benefit, until no further improvement in (5) is possible.

4. SIMULATION RESULTS

We consider two variants of our proposed service-centric approach: 1) A scheme that maximizes the overall user satisfaction and does not provide any guarantees for the less popular contents (Max-MOS+pop). 2) A scheme that defines a minimum guaranteed QoE for all upstreamed videos (Max-MOS-Fair+pop). We first run the Round Robin (RR) resource allocation scheme as a baseline to determine the minimum QoE a user should get. The minimum QoE of each user is then added as an additional constraint to solve the optimization problem in (5). Generally, an explicit guaranteed QoE value for each user or group of users can be defined (e.g., [14]). We also compare our service-centric approach with a QoE-driven one that does not consider video popularity (Max-MOS-Fair), and a reference RR scheme that allocates to each user an equal number of PRBs (RR).

We consider 25 different video sequences. At each simulation run, the assignment of the sequences to upstream users is shuffled to guarantee no particular sequence enjoys higher popularity. We set the Zipf shape parameter for video popularity to 1.0 [1]. The simulation parameters are summarized in Table 2.

We initially set the number of upstream users to 25, each uploading a different content. The number of downstream users is fixed to 1000. Figure 3 shows the cumulative distribution function (CDF) of the mean MOS for the different schemes. The mean MOS is computed by averaging the MOS for all downstream users over 200 simulation runs, 30 sec each. Please note that this is one way to measure the average user satisfaction while considering the downstream users' requests across different contents. The Max-MOS-Fair approach improves the mean MOS compared to the reference

Fig. 3. CDF of the mean MOS for 25 upstream users.

RR scheme. Both proposed service-centric approaches show an additional gain compared to the QoE-driven approach as they take the popularity of upstreamed contents into account. Meanwhile, the gain decreases in the Max-MOS-Fair+pop approach as a result of the constraint on minimum guaranteed QoE for the less popular contents.

The above results can be further explained by inspecting the distribution of utilities and resource shares for each content. Figure 4 shows the average resource share per content as a function of the rank of the content. Contents are indexed from 1 to 25 which represent the most and least popular contents, respectively. Both service-centric approaches allocate more resources for popular contents and the distribution of resource shares much reflects the Zipf popularity distribution. Again, the Max-MOS-Fair+pop scheme shows a less skewed distribution due to the fairness constraint. The other two schemes will allocate on average equal resources for each content as they do not consider the content popularity. Figure 5 shows the distribution of average utilities of the 25 uploaded contents, sorted by their popularity. The Max-MOS-Fair approach provides an average gain compared to RR for all contents, irrespective of the content popularity. The Max-MOS+pop scheme improves the utility of popular contents dramatically. Less popular contents (i.e., contents which receive fewer downstream requests) suffer a decrease in their individual utilities at the expense of higher overall user experience. The Max-MOS-Fair+pop scheme provides a slightly lower improvement in utility for popular contents compared to Max-MOS+pop, but it still guarantees a minimum QoE for the less popular contents (contents 13 to 25 get the same utility as in the RR scheme).

Figure 6 shows the average MOS experienced by the downstream users as we vary the number of upstream users. When the number of upstream users is low, there are enough uplink resources
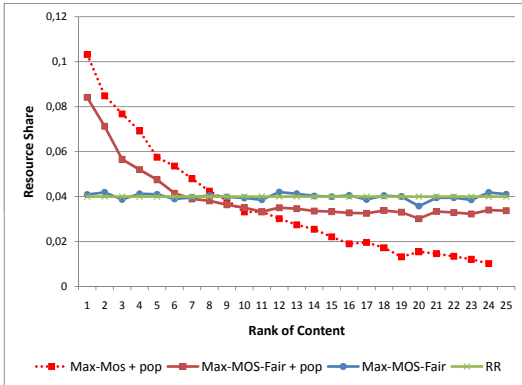
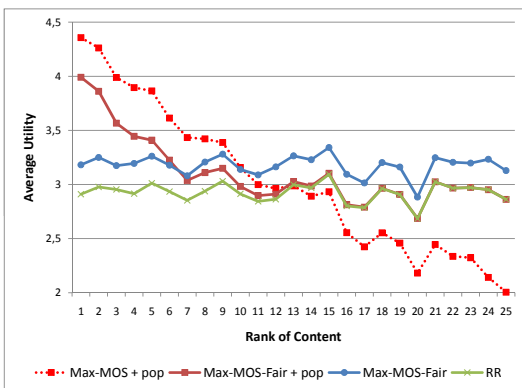**Fig. 4**. Average resource share for 25 upstream users.

Fig. 5. Average utility for 25 upstream users.

and the performance of the QoE-driven and the service-centric approaches is similar. As we increase the number of upstream users, the competition for resources is tighter and the service-centric approach improves the average MOS by prioritizing the popular contents. For a target MOS of 3.0, the Max-MOS-Fair+pop approach can admit 40 upstream users whereas the Max-MOS-Fair scheme can admit 30 users compared to 23 users for an RR scheme. This results in a system dimensioning gain of 25% and 42% compared to the QoE-driven and RR approaches respectively, while maintaining a minimum QoE for the less popular contents. The Max-MOS+pop scheme can even achieve a larger average gain, in particular as the network resources are more limited.

5. CONCLUSION

In this paper we introduced a service-centric approach that incorporates popularity feedback from a video portal into QoE-based uplink resource allocation. The objective is to improve the overall user satisfaction in loaded network situations by considering the asymmetry in the users' consumption of video contents. We addressed the fairness issue that could result from prioritizing the video contents and we set a minimum assured QoE for the less popular contents. The proposed approach is evaluated for the uplink of LTE and compared to a QoE-driven CLO scheme. We observed a significant and consistent gain by including the popularity information into the uplink resource allocation under various simulation scenarios.
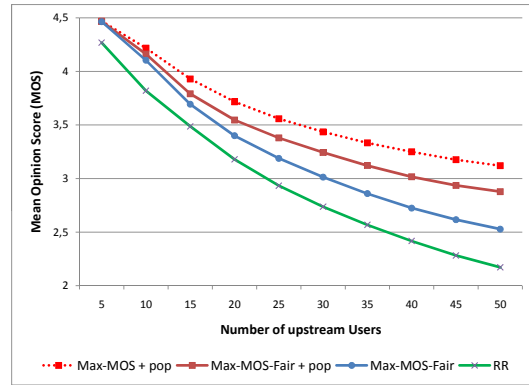

Fig. 6. Mean Opinion Scores averaged over 200 simulation runs.

6. REFERENCES

[1] M. Cha et al., "Analyzing the video popularity characteristics of large-scale user generated content systems," *IEEE/ACM Trans. on Netw.*, vol. 17, no. 5, pp. 1357–1370, Oct. 2009.

[2] P. Gill, M. Arlitt, Z. Li, and A. Mahanti, "YouTube traffic characterization: a view from the edge," *Proc. ACM IMC'07, San Diego, USA*, pp. 15–28, Oct. 2007.

[3] S. Sen et al., "Proxy prefix caching for multimedia streams," *Proc. IEEE INFOCOM'99, NY, USA*, March 1999.

[4] S. Borst et al., "An analytical model for provisioning of emerging personalized content distribution services," *Proc. ITC'09, Paris, France*, Sept. 2009.

[5] W. Rao, L. Chen, A. Fu, and G. Wang, "Optimal resource placement in structured peer-to-peer networks," *IEEE TPDS*, vol. 21, no. 7, pp. 1011–1026, July 2010.

[6] Fang Xie and Lei Du et al., "Popularity aware scheduling for network coding based content distribution in ad hoc networks," *Proc. IEEE PIMRC'07, Athens, Greece*, Sept. 2007.

[7] Y. Guo and H. Chaskar, "Class-based quality of service over air interfaces in 4G mobile networks," *IEEE Communications Magazine*, vol. 40, no. 3, pp. 132–137, March 2002.

[8] R. Berry and E Yeh, "Cross-layer wireless resource allocation," *IEEE Signal Processing Magazine*, vol. 21, no. 5, pp. 59–68, Sept. 2004.

[9] S. Khan and Y. Peng et al., "Application-driven cross-layer optimization for video streaming over wireless networks," *IEEE Communications Magazine*, vol. 44, no. 1, pp. 122–130, Jan. 2006.

[10] S. Thakolsri, S. Khan, E. Steinbach, and W. Kellerer, "QoE-driven cross-layer optimization for high speed downlink packet access," *Journal of Communications*, vol. 4, no. 9, pp. 669–680, Oct. 2009.

[11] L. Choi et al., "Sequence-level methods for distortion-rate behavior of compressed video," *Proc. IEEE ICIP'05, Genova, Italy*, Sept. 2005.

[12] M. E. J. Newman, "Power laws, Pareto distributions and Zipfs law," *Contemporary Physics*, vol. 46, pp. 323, 2005.

[13] M. Ivrlac and J. Nossek, "Cross layer design - an equivalence class approach," *Proc. IEEE ISSSE '04, Linz, Austria*, Aug. 2004.

[14] A. Saul, "Wireless resource allocation with perceived quality fairness," *Proc. IEEE ACSSC'08, CA, USA*, Nov. 2008.

[15] 3GPP TR 25.814 V.7.1.0, "Physical layer aspects for Evolved Universal Terrestrial Radio Access (E-UTRA)," Sept. 2006.

[16] 3GPP TR 36.942 V.8.1.0, "E-UTRA; Radio Frequency (RF) system scenarios," Dec. 2008.

[17] 3GPP TR 36.814 V.9.0.0, "E-UTRA; Further advancements for E-UTRA physical layer aspects," March 2010.

[18] D. Jurca et al., "Media flow rate allocation in multipath networks," *IEEE Trans. on Multimedia*, vol. 9, no. 6, pp. 1227 –1240, Oct. 2007.