Efficient Integration of Satellites in Collaborative Network Structures

Bernd Klasen

SES ASTRA TechCom, L-6815 Betzdorf, Luxembourg
bernd.klasen@ses.com

Abstract. Technological advances are enabling the reception of satellite broadcast even on mobile devices. However, mobile device usage is highly individual and exhibits an on-demand characteristic which rather suggests unicast. In order to avoid a wastage of the strengths of satellite broadcasts, additional effort is needed. This paper presents an approach to integrate satellites into heterogeneous collaborative networks for provision of IP services and broadband Internet access. Thus it supports both real time multimedia streaming and multimedia on demand.

Keywords: IP over satellite, network convergence, heterogeneous networks, content distribution.

1 Introduction

Since the establishment of the Internet and basically with the launch of the World Wide Web (WWW), human life considerably changed. It is affecting the way we work as well as the way we spend our leisure time. We look for Information, buy and sell, play games and communicate with other people in forums, chats and social networking platforms. The importance of online activities and the time we donate to them increases with the enhancing availability and growing bandwidth of the Internet accesses. At the same time we accustom ourselves to the way the Internet provides us with things we are looking for – which is immediate, ubiquitous and on–demand. Everything is just a few clicks away and waiting times of several seconds are considered as long in this context.

Due to the success of smart–phones and other mobile devices that encourage us to an ubiquitous Internet usage, this effect has even been further amplified. Our desire to exchange information, to participate in social networks, to use the latest software – the word apps is often used in this context – and to access videos and other multimedia files always and everywhere leads to constantly increasing network traffic and bandwidth demands. Permanent upgrades and investments on the infrastructure are needed in order to preserve connection speed and service availability under this growing load. Considering this situation, being now able to receive satellite broadcasts on mobile devices might seem to be the solution to all present problems. However, satellites – and broadcast networks in general – are inherently designed to serve a large number of people with the same content at the same time, which is in contrast to the highly

P. Pillai, R. Shorey, and E. Ferro (Eds.): PSATS 2012, LNICST 52, pp. 130–137, 2013.

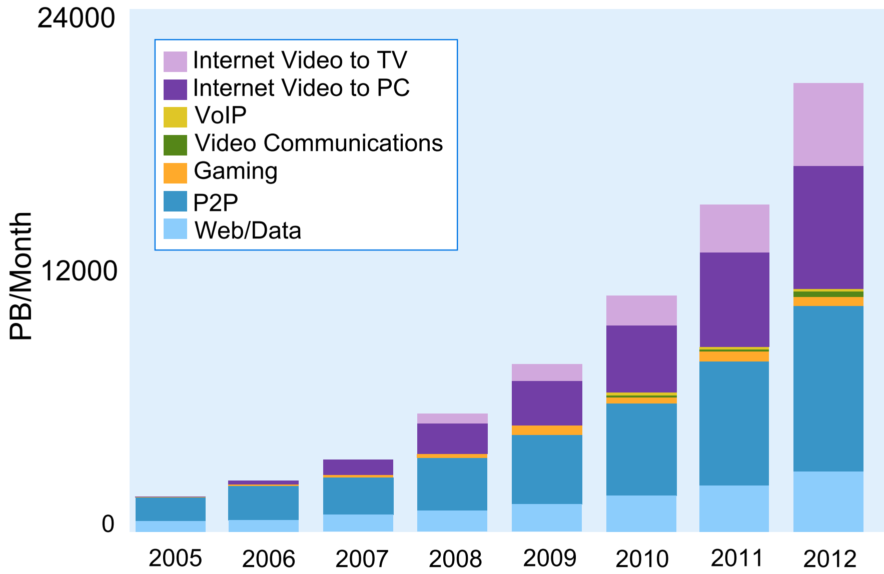

Fig. 1. The Internet traffic forecast. Source: www.cisco.com

individual behavior described above. In case the satellite bandwidth must be split among thousands or even millions of users for individual transfers, the possible benefits are extinguished. As a solution to this problem, this paper presents an approach for efficient satellite integration into mobile networks that uses resource request aggregation. Files that are downloaded by a large number of users nearly simultaneously are being broadcasted, while less popular data is send via unicast. The following sections will show that this approach not only will cut the peak loads caused by "hot" or "hype" content but also will allow better QoS for the most traffic intense type of online content. The latter, according to studies presented in [10] and to the Cisco Internet traffic forecast shown in figure 1, are video files. It is worth noting that although P2P traffic seems to cause at least the same amount of traffic, we have to recall that a large fraction of files exchanged via P2P consists of multimedia content. Due to the high relevance of video files for the overall Internet traffic, the proposed network model is motivated and described in the scope of video files. However, it can be also applied to other types of content with similar characteristics regarding the access patterns.

2 Related Work

The potential for traffic savings by aggregating viewers and then using an Internet multicast has been analyzed in [1]. While this study shows the applicability of request aggregation in an video on demand environment, it does not answer the question how the multicast can be realized. So far, multicasts over the Internet have not been successfully realized in a large scale due to the effort for multicast–group management. Related work has also been presented by [5] who concentrate on the server load reduction potential when using a P2P–like distribution model for video on demand. Further the authors of [4] analyzed also an P2P based video delivery model for YouTube which had a focus on the reception side. Also in [2] a P2P based approach is analyzed. None of these studies utilized a satellite or another broadcast network for increased efficiency. However, their research shows that a timely correlation of online video requests can be observed.

An approach focusing on the lower (packet) level of DVB service delivery to mobile clients and the corresponding reliability issues is presented in [7]. A satellite based content distribution model in hybrid networks has been presented in [8]. It relies on the distribution of files before the demand arises, which becomes possible due to popularity predictions, request forecasts, subscriptions and the availability of large caches at the users. However, the available memory on mobile devices is much more scarce than it is the case on TV–sets, Set–Top–Boxes or Internet modems, which can use cheap hard disks as caches that provide high capacities. Thus for mobile networks, a different approach must be taken. Let us assume that we have enough storage to completely cache the video currently played which is reasonable considering current standards.

3 Content Distribution Model

Since online video on demand (VoD) allows users to start playback of a video at an arbitrary moment in time, usually it is sent to the customer by the time playback is started as unicast, which is natural choice considering the individual nature of this service. But in case thousand people start watching at the same moment, the whole data is sent one thousand times. Besides of the massive amount of traffic this generates, there is another problem: If the average bandwidth of the user does not at least match the video's encoding bitrate, a fluent playback is not possible or he has to wait a long time before enough data is buffered. Considering the numerous rural areas where UMTS is still not available – not even to talk about LTE – this is an important aspect considering mobile video reception.

Before we start with the details of the model, let us be reminded on an important property that is essential for the applicability of the described approach. This property is the power law distribution of online video popularity which is confirmed by [6]. Similar results are presented by the authors of [3] who state that 10% of all available videos cause 90% of all web traffic. While many studies pay special attention on the long tail of that distribution, the model described in this paper concentrates on the opposite, on the most popular files. It has further been shown in [9] that there are online videos which exhibit heavy bursts in their request patterns. The fact that there exist extremely popular videos which at the same time show a strong timely correlation in their view statistics ensures that there are files which can efficiently distributed by using the model presented in this paper. As already mentioned, the latter is roughly based on the rule that files with many concurrent requests are broadcasted while the others are transmitted via the terrestrial network as unicast. In detail, the model works as follows:

Like it is common for peer-to-peer file exchange protocols, we subdivide each file into several pieces of size S_P. We choose the piece size accordingly to the available mobile network bandwidth B_M so that one piece can be submitted per second. This also means that this is the maximum encoding bitrate that can be used for the video in order to allow fluent playback. Later we will also see simulation results where a higher video encoding bitrate is used which means that these videos could not have been viewed instantaneously in full quality when only the mobile network with EDGE bandwidth was used for delivery. Regardless of a specific encoding bitrate, given a size S_{F_i} for a file F_i, this means that F_i is split into

$$|P_{F_i}| := \frac{S_{F_i}}{S_P} \quad (1)$$

pieces. Let further the available satellite bandwidth be B_S. Whenever the number of active downloaders exceeds a defined broadcast threshold BT, a broadcast is scheduled. The full satellite bandwidth is utilized then, which means that we are able to transmit multiple file pieces within one second. The ratio r between the files transmitted per second via mobile network and via broadcast depends on the available mobile network bandwidth. Thus it is

$$r = \frac{B_S}{B_M} \quad (2)$$

Since we assume that all pieces are received sequentially via mobile network, when BT is reached we check which is the piece with the highest sequence number $piece_{max}(F_i)$ that at least one on the active downloaders already holds. Then during the next step respectively the next second, r pieces are submitted at once via broadcast. In case that

$$piece_{max}(F_i) + r > |C_{F_i}| \quad (3)$$

we change the start sequence number for the broadcasted pieces to $|C_{F_i}| - r$ in order to utilize the full available bandwidth. This means in that special case some pieces are broadcasted even if BT is not reached for them, but it further increases the download speed for the remaining clients. Also, all clients that would not have reached the sequence number $piece_{max}(F_i)$ after the broadcast, the next piece in line is transmitted via unicast just as it would have been the case without a broadcast. Two other parameters we need are the maximum number of clients that can start watching a video per second (CpS) and the probability P that they will. The following evaluation of this model will show the possible traffic savings.

4 Evaluation

The first important question concerning the evaluation of the proposed model is the choice of the parameters. Regarding the mobile network bandwidth, choose EDGE with 236 Kbit/sec as the lowest common denominator since it is still the maximum available network bandwidth for several regions.Thus let

$$S_P = 236kbit \tag{4a}$$

$$B_M = 236kbit/sec \tag{4b}$$

$$B_S = 36,000kbit/sec \tag{4c}$$

A very critical parameter is the number of users that might request a certain file per step respectively per second. Here we use results of former studies that analyzed video request rates on YouTube [9]. According to those, there are videos which exhibit an average of 715.5 requests per minute during the first three days after they have been uploaded.

This means we have approximately 12 requests per second. Since the data is relatively coarse grained, it is impossible to distinguish how these requests are distributed within a period shorter than one hour. Thus we assume a Poisson distribution and let the

$$CpS = 7,155 \tag{5a}$$

$$P = 0.0017 \tag{5b}$$

This corresponds to the probability for a client to request the file within 10 minutes. Further we use the average video duration of 342 seconds, which relies on the results for very popular videos from the YouTube study mentioned above. Considering the now commonly used average bitrate of $1,057Kbps$ for

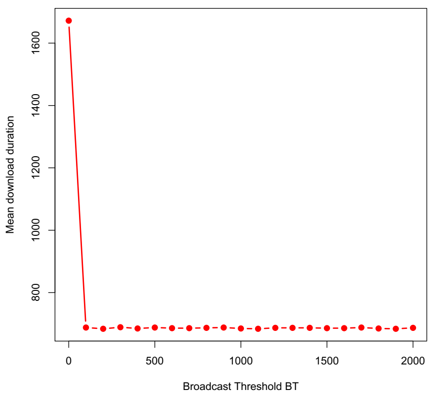

Fig. 2. Encoding bitrate = 1057 Kbps

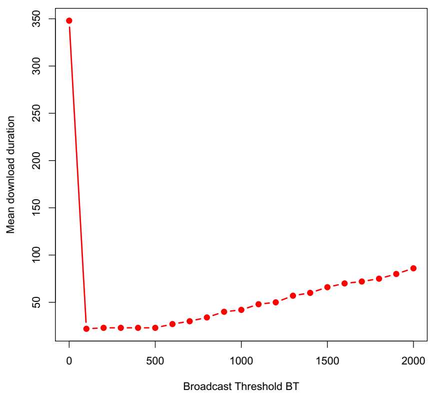

Fig. 3. Encoding bitrate = 236 Kbps

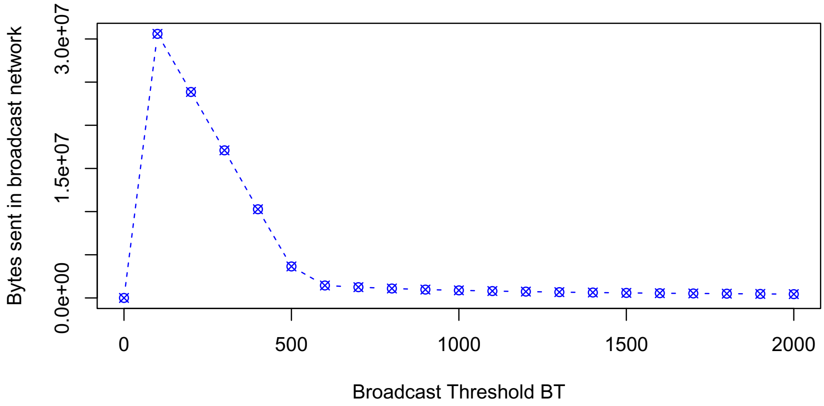

Fig. 4. kByte broadcasted

video with resolution 480p – which is not the highest quality but a common resolution we find on screens of mobile devices – and a audio bitrate of $96Kbps$ it results in an average video file size of $\approx 48.16MByte$. We also use a lower bitrate of $236Kbps$ with a corresponding filesize of 10.01 MB which allows a fluent playback without an excessively long buffering period when we have no broadcasts. Figure 2 and 3 show the download durations for varying values of BT for both encoding bitrates, while $BT = 0$ means that broadcasts are disabled. For the higher bitrate (figure 2) we have an average bandwidth saving in the mobile network of 58%, while it is more than 86% in case of the lower bitrate (figure 3). These are average values, and what we can clearly see in figure 3 is that the time a client needs to complete the download immediately increases

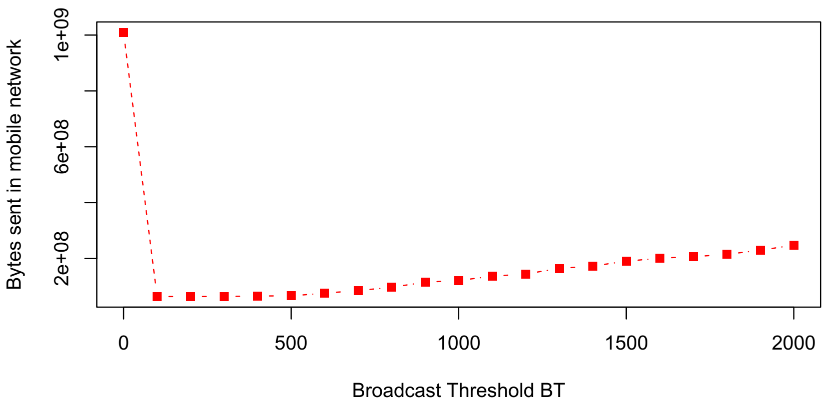

Fig. 5. kByte unicasted

with growing values of BT. This is what one expects considering the smaller filesize. In general it can be stated that the larger a file – and thus the longer the expected download time – the higher we can choose BT while still achieving a high mobile network traffic reduction. Figure 4 shows the broadcast bandwidth, figure 5 the mobile bandwidth needed under changing values of BT. Again, for $BT = 0$ the broadcast is disabled.

5 Conclusion

This paper presents a content distribution model that uses satellite broadcasts in a collaborative manner together with a terrestrial unicast network in order to reduce network traffic. It has been shown that the unicast traffic can be reduced by more than about 86%. At the same time, only relatively little satellite bandwidth is needed. This approach requires that some authority – for example the mobile or satellite network provider – keeps track of the number of current requests for a certain file and decides on the delivery channel accordingly. Since this violates the net–neutrality, it is important to implement the described approach in a way that users are aware of this and use it voluntarily. The reason why people should freely decide to do this is a fluent video playback even under high network load and shorter download times for videos. The latter can be very important if people spend also time in environments where no network access is possible.

Acknowledgment. This research is supported by the *National Research Fund* (FNR) of Luxembourg.

References

1. Aggarwal, V., Caldebank, R., Gopalakrishnan, V., Jana, R., Ramakrishnan, K.K., Yu, F.: The effectiveness of intelligent scheduling for multicast video-on-demand. In: Proceedings of the Seventeen ACM International Conference on Multimedia - MM 2009, p. 421 (2009)
2. Carlsson, N., Eager, D., Mahanti, A.: Peer-assisted on-demand video streaming with selfish peers. In: Networking 2009, pp. 586–599 (May 2009)
3. Cha, M., Kwak, H., Rodriguez, P., Ahn, Y.-Y., Moon, S.: I tube, you tube, everybody tubes: analyzing the world's largest user generated content video system. In: IMC 2007: Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement, pp. 1–14. ACM, New York (2007)
4. Cheng, X., Liu, J., Wang, H.: Accelerating YouTube with video correlation. In: Proceedings of the First SIGMM Workshop on Social Media - WSM 2009, p. 49 (2009)
5. Garbacki, P., Epema, D.H.J., Pouwelse, J., Van Steen, M.: Offloading servers with collaborative video on demand. In: Proceedings of the 7th International Conference on Peer-to-Peer Systems, p. 6. USENIX Association (2008)
6. Gill, P., Arlitt, M., Li, Z., Mahanti, A.: Youtube traffic characterization: a view from the edge. In: Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement, pp. 15–28. ACM (2007)

7. Gotta, A., Barsocchi, P.: Experimental video broadcasting in DVB-RCS/S2 with land mobile satellite channel: a reliability issue. In: Fourth International Workshop on Satellite and Space Communications 2008 (2008)
8. Klasen, B.: Efficient Content Distribution in Social Aware Hybrid Networks. Journal of Computational Science (2011)
9. Klasen, B.: Social, fast, efficient: Content distribution in hybrid networks. In: 2011 IEEE Symposium on Computers and Communications, ISCC, pp. 61–67. IEEE, Kerkyra (2011)
10. Maier, G., Feldmann, A., Paxson, V., Allman, M.: On dominant characteristics of residential broadband internet traffic. In: Proceedings of the 9th ACM SIGCOMM Conference on Internet Measurement Conference, vol. 9, pp. 90–102. Citeseer, New York (2009)