Energy Consumption Analysis of Peer-assisted Video-on-Demand Streaming in Managed Networks

Sasho Gramatikov

Faculty of Computer Science and Engineering, University "Ss. Cyril and Methodius", Skopje, Macedonia
sasho.gramatikov@finki.ukim.mk

Abstract. The delivery of Video-On-Demand (VoD) contents is one of the leading services that occupies most of the Internet traffic with tendency to dominate in the years that come. Although it offers a great convenience to the users, it is a serious burden for the privately managed networks which have to process huge amounts of traffic meeting high levels of quality of service. This high traffic-demanding service not only requires higher number of network devices and streaming servers, but also demands high energy consumption which is costly for the network operator and harmful for the environment. Therefore, in this paper, we analyse the energy consumption of a VoD delivery system where the clients partially participate in the streaming and thus significantly reduce the overall traffic in the core of the privately managed network. We show that the the peer-assisted VoD streaming significantly makes the VoD delivery greener service since it reduces the energy consumption both by the network operator and the clients.

Keywords: Energy consumption, Video-on-Demand, P2P, Streaming

1 Introduction

The advances of network technologies and the variety of offered video services have made the videos one of the most popular contents on the Internet. Although they are very traffic demanding due to their large size, they have become easily accessible and shared among a growing community of users. Therefore, the video traffic has dominated the global traffic in the Internet nowadays and has a growing tendency to take even more dominant position occupying 80-90% of the globally exchanged traffic worldwide [4]. A significant part of this traffic will belong to the Video on Demand service (VoD) which will triple the amount of traffic that it generates nowadays. The reason for this rapid growth is the increasing number of video breaking the boundaries of the choice of contents and the time they are watching. This convenience is, however, a great concern for the operators since every video requires a separate unicast stream causing congestion in the core of the network with the growing number of users and the demand for higher quality videos.

In order to maintain the required quality of service, the operator has to install additional network equipment, streaming servers and higher capacity links. Apart of the increased installation costs, the operator is faced with more concerning issue, which is the increased energy consumption. The energy consumption for delivering digital contents in general is a topic that concerns many scientist in the field of information and communication technologies since this sector is responsible for nearly 5% percent of the total energy consumption of the developed countries [10], out of which 1% belongs to the Internet traffic. Half of this energy is consumed for switching and transmission of the data [6]. Therefore, modelling and estimating the energy consumption of a video delivery systems is an important task towards building more energy efficient solutions. In [1] the authors measure the energy consumption of various network devices and gigabit network cards [2]. An overall estimation of energy consumption per data bit in the Internet, taking into account various access technologies is given in [10][5]. In [7], an analysis of energy consumption for switching, transporting, processing and storing data in cloud computing is presented. In [11], the authors estimate the energy consumption for VoD delivery in privately managed network. The authors also propose energy efficient algorithm for content placement in a hierarchy of servers in order to reduce the overall energy consumption.

In this paper, we analyse the energy consumption for peer-assisted Video-on-Demand streaming in privately managed networks [9]. The main goal of the system is to concentrate the streaming traffic close to the clients, to alleviate the streaming servers and reduce the traffic in the core of the network. It achieves these goals by implementing algorithm for optimal placement of the contents in the network based on the clients' activity and state of the network [8], and by taking the advantage of unused uplink capacity of the peers to stream pre-stored contents to other peers in the community. Since this approach reduces and optimizes the traffic in the network, the focus of this paper is the energy consumption for streaming and transmission of the video traffic to the clients from the streaming servers to the clients' TV sets. We show that the peer-assisted streaming can significantly reduce the power consumption in the distribution network of the service provider on account of the streaming resources offered by the clients. In order to further decrease the overall energy consumption, we also analyse the energy savings that can be achieved if the clients turn off their Set-Top Box (STB) upon the end of each video session. Although this behaviour reduces the available streaming resources for the operator, and hence increases the energy consumption, the overall energy consumption per client will be considerably reduced.

The rest of the paper is organized as follows: in Section 2 we describe the system for peer-assisted VoD streaming and the network architecture used in the analysis. After we present the method for calculation of the energy consumption for video delivery in 3, we describe the simulation scenarios and analyse the results obtained in Section 4. Eventually, we give the conclusions from our work in Section 5.

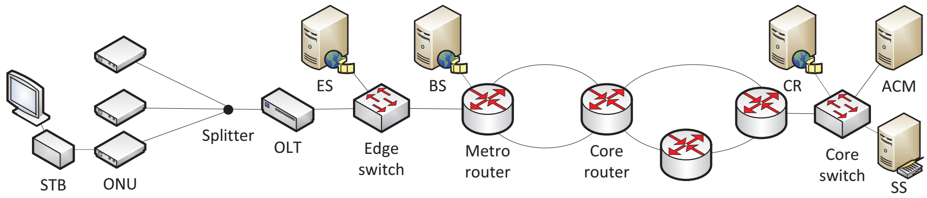

Fig. 1. Network architecture

2 System description

The system for opeer-assisted VoD streaming is designed for private networks owned and managed by a company that offers VoD streaming and other TV and data related services. The provider has the overall control of the network and can dimension it according to the number of subscribed users. The system consists of streaming and management servers placed in different locations of the core, metro and access network. The access network technology used in this work is Passive Optical Network (PON) which offers Fiber to the Home (FttH) service to the clients, i.e., each clients is connected to the provider with high-speed optical link. This technology requires that the provider installs Optical Line Terminal (OLT) placed in the end-offices. Each port of the OLT is connected to a splitter, a passive component that broadcast the same signal to up to 64 clients. On the other side of the optical link, the signal is handled by the Optical Network Unit (ONU) which is a device on the customers' premises. Each client posses a Set-Top Box (STB) with Digital Video Record (DVR) functionality, enabling storage of certain amount of video contents. Since the provider has control over the network, it assigns certain part of the uplink and storage capacity of the STBs to place portion of the offered videos for the purpose of the peer-assisted streaming process. The clients connected to the same OLT form a local community with size n. One local community is serviced by a streaming server, called Edge Server (ES). The ES has limited storage and streaming capacity, and therefore, hosts only the most frequently requested video contents.

The OLT and the ES are directly connected to the edge switch, which is then connected to the metro router of the metro network. The provider places streaming servers, called Branch servers (BS), in the Video Hub Offices (VHO) of the metro networks to handle the requests for the less popular videos. The core network connects the routers of the metro networks by high capacity optical links. At the top level of this network hierarchy, in the Super Head End (SHE), resides the central repository server (CR) that hosts the entire video library. This server handles all the request for videos that cannot be served by the servers closer to the clients. Other important components of the system are the management servers represented by the Automatic Content Movement (ACM) server and the Service Selection (SS) server.

The ACM server communicates with the streaming servers, monitors their state and takes redistribution decisions. Whenever it detects that there are overloaded servers or that the users' behaviour has changed, it runs an algorithm for content redistribution. Using popularity data of the contents in the recent past, previously obtained from the SS server, the algorithm decides whether a replica of a content item should be moved to another server, removed or left as it is. The execution of the algorithm results with a new distribution of the video contents, which is deployed by execution of a set of removal and replication commands issued by the ACM. The SS server is responsible for accepting the clients requests and redirecting them to the most appropriate streaming server. It implements a redirection strategy that uses the current state of the system and the availability of the contents to take a decision on where to redirect the clients.

The peer-assisted streaming is implemented by taking advantage of the unused streaming and storage capacity of the clients. Since the uplink capacity of the clients is lower than the streaming rate r of the High Definition (HD) videos offered by the operator, the contents are divided into m parallel strips. Thus, the streaming rate of each strip, also called a channel, is m times smaller than the real streaming rate of the video, permitting the peers to partially participate in the streaming process. Depending on the streaming capacity of the peers, expressed as number of channels k, the uninterrupted streaming is provided by parallel streaming of m strips from different peers and their assembly at the receiving STB. Apart from storing the entire library of C video contents in the streaming servers, the provider stores strips of the contents in the STBs according to popularity distribution scheme. The storage capacity of the STBs reserved by the operator for the peer-assisted streaming is s contents.

The process of requesting video contents is modelled as Poisson process with average rate $\lambda = 1/w$, where w is the average time the client spends in idle state waiting to make a request. The process of serving the videos, i.e., the average duration of watching session is also considered as a Poisson process with average rate $\mu = 1/d$, where d is the average duration of the video session. Each request is first addressed to the ES, which keeps track of the availability of the contents on the peers in the same local community and their available streaming capacity. The ES finds the available peers that contain strips of the requested content and responds to the requesting client with the location of the assigned peers to serve the available strip. Afterwords, the requesting peer establishes connections with the assigned peers which deliver the strips they host. If not all the requested strips can be found on the peers, the requesting peer forwards its request to the SS which assigns the most appropriate streaming server to stream the missing strips.

In order to reduce the power consumption of the STBs, the clients can turn off their STB after they finish watching the requested video with probability p_{off} and stay in that state until the next request. Same as the process of requesting and serving videos, the process of failures is also modelled as a Poisson process with the same average rate as the requesting process $\lambda = 1/w$, i.e., the time the client spends waiting to make request with the STB turned on is the same as the time it waits to turn on the STB and make a new request for a video. With such a definition of the model, a client can be in one of the following states: waiting state, receiving state and off-state. The clients that are in the waiting or the receiving state are also called active clients since they can participate in the peer-assisted streaming.

Despite the energy savings, this behaviour yields two disadvantages from operators point of view: the clients that turn off their STBs cannot participate in the peer-assisted streaming, and the interruption of the streams that were currently streamed by these peers requires additional resources. The solution for providing uninterrupted streaming experience is delivery of the remaining contents by the streaming servers. The effect of these disadvantages is the increased traffic in the core of the network originating from the streaming servers, and consequently, increased energy consumption.

3 Energy consumption for video delivery

The most intuitive way to calculate the energy consumption in the peer-assisted VoD streaming system is to sum the energy consumed per streamed bit and the energy consumed for transmitting each bit through the network elements and then multiply this value by the total amount of traffic. One of the main issues for using this approach is the calculation of energy consumption per bit since the network devices consume energy even when they are in idle mode. According to the measurements of various network devices under different load in [1], the energy consumption in the range between idle mode and the maximum load is proportional to the load. However the coefficient of proportionality varies for different devices and in many cases turns out that the difference between the idle and maximum power state is minor and the amount of traffic handled by the network device has insignificant influence on the energy consumption. The same observation is taken when measuring the power consumption of Network Interface Cards (NIC) in [2].

Although apparently not accurate, the most common approach for measuring the energy consumption in large-scale networks is considering that the energy consumption is proportional to the throughput. In a large-scale system with traffic demanding services, the size of the network depends on the number of customers subscribed for the service. The more customers incorporate in the network, the more load has to be processed by the network devices. When these devices approach the threshold of maximum utilization, the network operator installs new network devices. Therefore, under assumption that almost the entire capacity of the network devices is utilized for processing the requested video traffic, the assumption that the energy consumption of the devices is proportional to the throughput will lead to satisfactory results.

In our analysis, we will measure the energy consumption P_{cl} for receiving and streaming videos by the clients and the energy consumption for delivering the videos by the provider P_{pr}. The energy consumption on the clients' premises will be calculated as:

$$P_{cl} = P_{ONU} + p_{act} P_{STB} \tag{1}$$

S. Loshkovska, S. Koceski (Editors): ICT Innovations 2015, Web Proceedings, ISSN 1857-7288

where P_{ONU} and P_{STB} are power consumed by the ONU and the STB, accordingly, and p_{act} is the probability that a STB will be active. In the equation, we deliberately omit the energy consumption of the TV display since the clients can use a variety of TV sets with different energy specifications. One can also note that this equation does not depend on the amount of traffic since the ONU is always on and the STB is considered to consume the same energy amount for receiving and streaming videos. The probability p_{act} can be calculated by dividing the number of active clients by the total number of clients in the community. In our previous work [9], we calculate the average number of active peers as:

$$n_{act} = \frac{1 - p_{off} + \lambda/\mu}{(1 + p_{off} T_{off} \mu)\, \lambda/\mu + 1 - p_{off}} n \tag{2}$$

where T_{off} is the average time the client keeps its STB turned off. Since in this work we assume that the client spends the same time in the off-state as in the idle state, by substituting $T_{off} = w = 1/\lambda$ in (2) and dividing the equation by n, we can calculate the probability that a client is active p_{act} as:

$$p_{act} = 1 - \frac{p_{off}}{\lambda/\mu + 1} \tag{3}$$

The video delivery power consumption per customer is calculated by measuring the throughput of the network elements and the streaming servers' load and using the assumption that the power consumption is proportional to the applied load for large-scale system. The power consumption largely depends on the network configuration. In our work, we are considering a network as shown in Figure 1. Hence, the P_{op} will be expressed as

$$\begin{aligned} P_{pr} = \frac{r}{m}\frac{2}{N}((U_{ES} + U_{BS} + U_{CR})(P_{OLT} + P_{edge} + P_{srv}) + \\ + (U_{BS} + U_{CR})P_{metro} + U_{CR}(2P_{core} + P_{switch}) + U_{p2p}P_{OLT}) \end{aligned} \tag{4}$$

where U_{ES}, U_{BS}, U_{CS} and U_{p2p} are the number of strips streamed by ES, BS, CR and the peers, accordingly. The first term in the equation refers to the traffic that passes through the access network, i.e., through the OLT and the edge switch. This traffic originates from the streaming servers and on its way to the clients has to be handled by the edge switch and the OLT. The second term refers to the streaming traffic from the BS and CR that passes through the metro network routers. The third term is the traffic from the CR placed in the SHE, processed by the core switch and core routers. Since the network has one SHE, the distance from the the metro networks to the SHE may be different. In our work, we are considering a network where the mean number of routers that will process this traffic is 2, although for larger systems this number is well above this value [11]. The last term of the equation is the P2P streaming traffic that originates from the clients of the local community and terminates in the same local community. We multiply the entire expression by factor of 2 since for every spent Watt of a device, there is additional Watt consumed for cooling the device [11]. The equation is also multiplied by the streaming rate of each strip r/m and divided by the total number of clients in the system in order to obtain the energy consumption per client.

4 Simulations and results

In order to measure the energy consumption for peer-assisted video streaming, we developed a simulation model in Matlab. In the simulations, the network consists of $N = 100000$ clients groped in local communities of $n = 1000$. The operator offers a library of 1500 HD movies with playback rate $r = 10$ Mbps. The videos are divided into $m = 10$ strips and stored in both the streaming servers and the STBs of the peers. The streaming rate of each channel is $r/m = 1$ Mbps. The ES, BS and CR have storage capacity to host 20%, 40% and 100% of the contents, accordingly and streaming capacity to handle 20%, 33% and 100% of the total number of simultaneous streams. The peers have STB with storage capacity of $s = 10$. The clients are generating requests for videos with average waiting time $w = 1/\lambda = 40$ min and watch the requested video in average $d = 1/\mu = 80$ min.

The videos are previously distributed in the peers according to a distribution scheme that dedicates 30% of the storage space of the STBs for uniform distribution of the popular videos and the rest of the storage for uniform distribution of the not popular contents. According to the 20-80 rule [13], the first 20% of the most requested videos in the library are considered as popular, while the rest are not popular.

The popularity of the video contents obeys the Zipf-Mandelbrot (ZM) distribution [12] with skew factor $\alpha = 0.8$ and shifting constant $q = 10$.

Based on the measurements of the energy consumption of various network devices in [10] [11][3], we will be using the values presented in Table 1.

P_{core}	P_{metro}	P_{edge}	P_{OLT}	P_{serv}	P_{ONU}	P_{STB}
15 nJ/b	15 nJ/b	28 nJ/b	19 nJ/b	20 nJ/b	5 W	30 W

Table 1. Overview of the enery concumption of equipement used in the simulations.

In our first simulation scenario, we measure the average energy consumed by the operator for peer-assisted streaming of VoD contents for various streaming capacities of the peers. As it is expected, the higher streaming capacity of the peers implies higher participation of the peers in the streaming process which alleviates the streaming servers and reduces the traffic in the core of the network. The result of such reduction of the traffic is also reduction of the energy consumption for delivery of the videos by the provider. One must note that the clients still consume the same energy levels since we assumed that the STB has a constant energy consumption no matter whether the clients watch a video or not. The results of the measurements of the energy consumption per client as a function of the streaming capacity of the peers are shown in Figure 2. The figure shows that the increment of the number of channels linearly reduces the energy consumption up to the capacity of 4 channels, i.e., 4 Mbps, where the energy consumption per user is reduced 50% compared to the case with pure server streaming. The importance of this reduction can be realized if we calculate the energy consumption savings of large-scale system with more than 10^6, which according to the results from the figure would be 600 kW. For higher streaming capacities, the further increment of the streaming capacity of the peers has insignificant effect on the energy consumption because not all the requested strips are available on the peers and there is still energy consumption by the OLT which has to handle all P2P traffic among the peers.

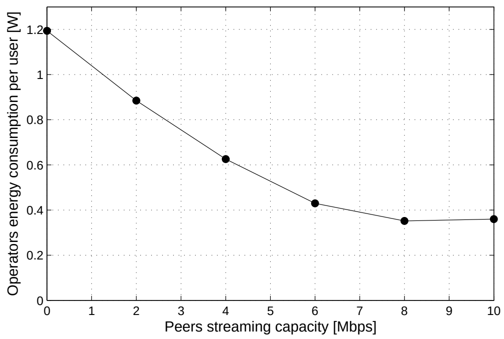

Fig. 2. Dependency of the operators power consumption per client on the size of the local community and the peers' streamng capacity

In the next simulation scenario, we measure the energy consumption on the clients' and provider's premises obtained by letting the clients turn off their STB with probability p_{off}. In Figure 3, we show the energy consumption by the clients obtained from (1) and 3. The energy consumption decreases linearly with the probability p_{off} with steepness that depends on the ratio λ/μ. This ratio defines the activity of the clients and their utilization of the streaming services. The more active the clients are, the more they participate in the streaming and the more energy is consumed. Therefore, even in the cases when the clients turn off their STB with certainty, the power does not reach the minimum value of energy consumption of the ONU. The maximum power savings that can be reached by the clients are 10 W, which is a significant reduction in our case of large-scale system.

The energy saving on the cleints' premises, however, comes with a drawback related to the energy consumption of the provider. The increased number of clients which turn off their STB reduces the available resources for peer-assisted streaming and increases the number of interrupted streams that have to be

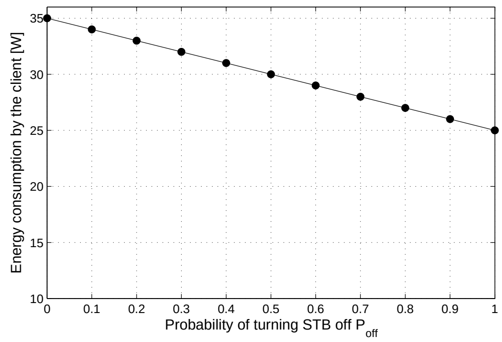

Fig. 3. Dependency of the operators power consumption per client on the size of the local community and the peers' streamng capacity

recovered by the streaming servers. Therefore, it is expected that the amounts of traffic originating from the streaming servers increases. The results of this behaviour for various streaming capacities of the clients are shown in Figure 4. From the figure, we can see that the increment of energy consumption increases linearly with the probability p_{off}. The maximum energy savings that can be achieved depend on the streaming capacity of the clients, however, if we compare any of the curves of the energy consumption per client by the operator in Figure 4 with the energy consumption by the clients in Figure 3, we can figure out that there is an order of magnitude more energy saving by the clients, than increased energy consumption by the operator. Therefore, we can conclude that by stimulating the clients to turn off their STB, the peer-assisted streaming becomes more energy efficient solution for delivery of VoD contents.

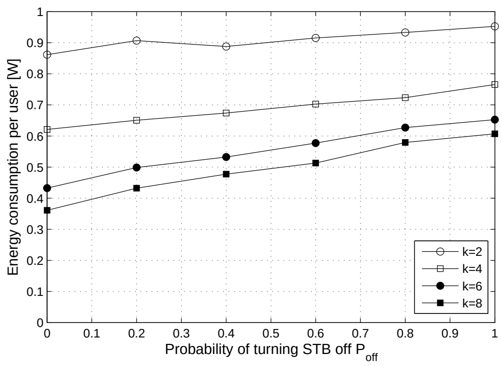

Fig. 4. Dependency of the operators power consumption per client on the size of the local community and the peers' streamng capacity

5 Conclusion

In this paper we present an energy consumption analysis of a large-scale system of peer-assisted VoD streaming in managed networks. The system takes advantage of the unused streaming and storage capacity of the peers to alleviate the streaming servers and to reduce the traffic in the core of the network. In our work, we measure the amounts of traffic originating from the streaming servers and processed by the network devices, and by using typical values of energy consumption of all the devices included in the peer-assisted streaming, we show that there are significant energy savings for video streaming by the VoD service provider. We also show that the global energy consumption can be further reduced if the clients turn off their STBs after they finish watching a video. Although this behaviour slightly increases the power demands by the network equipment of the operator, it proves to be ore energy efficient and greener solution for delivery of VoD contents.

Acknowledgement

This work was partially financed by the Faculty of Computer Science and Engineering at the "Ss. Cyril and Methodius" University.

References

1. A power benchmarking framework for network devices. Lecture Notes in Computer Science 5550 LNCS, 795–808 (2009)
2. Characterizing 10 Gbps network interface energy consumption. Proceedings - Conference on Local Computer Networks, LCN pp. 268–271 (2010)
3. Annual Report: Voluntary Agreement for Ongoing Improvement to the Energy Efficiency of Set-Top Boxes. Tech. rep., D&R International (2013)
4. Cisco Visual Networking Index: Forecast and Methodology, 2012-2017. Tech. rep., Cisco Systems (2013)
5. Baliga, J., Ayre, R., Hinton, K., Sorin, W., Tucker, R.: Energy Consumption in Optical IP Networks. Journal of Lightwave Technology 27(13), 2391–2403 (Jul 2009)
6. Baliga, J., Ayre, R., Hinton, K., Tucker, R.: Energy consumption in wired and wireless access networks. IEEE Communications Magazine 49(June), 70–77 (2011)
7. Baliga, J., Ayre, R.W.a., Hinton, K., Tucker, R.S.: Green Cloud Computing: Balancing Energy in Processing, Storage and Transport. Proceedings of the IEEE (2010)
8. Gramatikov, S., Jaureguizar, F., Cabrera, J., Garcia, N.: Content Delivery System for Optimal VoD Streaming. In: Proceedings of IEEE ConTEL. pp. 487–494 (june 2011)
9. Gramatikov, S., Jaureguizar, F.: Modelling and analysis of non-cooperative peer-assisted vod streaming in managed networks. Multimedia Tools and Applications pp. 1–28 (2015)
10. Hinton, K., Baliga, J., Feng, M., Ayre, R., Tucker, R.S.: Energy Efficiency in the Internet (April), 6–12 (2011)

11. Jayasundara, C., Nirmalathas, A., Wong, E., Chan, C.A.: Improving Energy Efficiency of Video on Demand Services. Journal of Optical Communications and Networking 3(11), 870 (2011)
12. Krogfoss, B., Sofman, L., Agrawal, A.: Caching Architectures and Optimization Strategies for IPTV Networks. Bell Labs Technical Journal, 13(3), pp. 13–28 (2008)
13. Yu, H., Zheng, D., Zhao, B.Y., Zheng, W.: Understanding User Behavior in Large-scale Video-on-Demand Systems. ACM SIGOPS 40(4), pp. 333–344 (Apr 2006)