



Reducing Energy Consumption in IPTV Networks by Selective Pre-Joining of Channels

Fernando M.V. Ramos
Computer Laboratory and
Department of Engineering
University of Cambridge
Cambridge, United Kingdom
fernando.amos@cl.cam.ac.uk

Richard J. Gibbens
Computer Laboratory
University of Cambridge
Cambridge, United Kingdom
richard.gibbens@cl.cam.ac.uk

Fei Song
Beijing Jiaotong University
Beijing, China
song2000fei@gmail.com

Pablo Rodriguez
Telefonica Research
Barcelona, Spain
pablo.rodriguez@tid.es

Jon Crowcroft
Computer Laboratory
University of Cambridge
Cambridge, United Kingdom
jon.crowcroft@cl.cam.ac.uk

Ian H. White
Department of Engineering
University of Cambridge
Cambridge, United Kingdom
ihw3@cam.ac.uk

ABSTRACT

IPTV services are the fastest growing television services in the world today. This is a bandwidth intensive service, requiring low latency and tight control of jitter. To guarantee the quality of service required, service providers opt to multicast all TV channels at all times to everywhere. However, a significant number of channels are rarely watched, so this method is provably resource- and energy-inefficient.

In this paper, we argue that the expected increase in quantity and quality of TV channels will become a serious issue, both in terms of bandwidth and energy costs. To overcome this problem, we propose a dynamic scheme that pre-joins only a selection of TV channels. This scheme was evaluated by means of trace-driven simulations using a large dataset from a commercial nationwide IPTV service. The dataset comprises 255 thousand users, 150 TV channels, and covers a 6-month period.

We show that by using our scheme IPTV service providers can save a considerable amount of bandwidth while affecting only a very small number of TV channel switching requests. To understand how these bandwidth savings are translated in energy savings, we developed a power consumption model for network equipment based on real measurements. The main conclusions are that while today the bandwidth savings will have reduced impact in energy consumption in the core network, with the introduction of very high definition channels this impact will become significant, justifying the use of resource-efficient distribution schemes such as the one proposed.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Green Networking 2010, August 30, 2010, New Delhi, India.
Copyright 2010 ACM 978-1-4503-0196-1/10/08 ...\$10.00.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Network Communications; C.2.3 [Network Operations]: Network Management

General Terms

Performance

Keywords

IPTV, energy friendly networks, green networks

1. INTRODUCTION AND BACKGROUND

We have been depleting the natural environment since the times of the industrial revolution. There is a scientific consensus that our planet will be unable to provide long-term support if this trend persists. Today, the Internet (excluding home networks, PCs and datacenters) consumes about 0.5% of the current electricity supply of a typical OECD nation. Although this still represents a relatively small share of the global energy consumption, this fraction is expected to increase quickly[2, 21].

Global IP traffic will quintuple from 2008 to 2013, at a compound annual growth rate of 40%[10]. The sum of all forms of video (IPTV, VoD, P2P) will account for over 91% of global consumer traffic by 2013. Due to the growing importance of this type of services, we have decided to focus on IPTV. In particular, we target broadcast IPTV, rather than Video on Demand. This is a resource intensive service, requiring high bandwidth, low latency and low jitter. Each video stream is encoded at a bit rate that can vary from around 4Mbps (SDTV) to 20 Mbps (HDTV). In the future this figure may increase by one or two orders of magnitude, with the advent of ultra high definition video standards (2K, 4K, UHDV)[12]. IPTV networks will thus need to be made more resource- and power-efficient than they are today.

Power consumption at the still low bit access rates of today is dominated by the access network. At high access rates, however, the energy bottleneck will move to the core[4]. Electronic routers will dominate energy consumption[3], while the power consumed by links will represent

only a small percentage of the total[21]. For this reason, in this work we decided to focus on the routing equipment.

Networks are provisioned at present for the worst case scenario and many times overprovisioned. In normal operation they are thus underutilised. There is, therefore, room for energy savings. The problem is that nowadays energy cost is a function of capacity, not throughput. For this reason energy proportionality [5] is becoming a primary design goal in network systems: energy consumption should be proportional to the amount of work performed. This has been a topic of research in wireless networks [15] and computer architectures[19] for some years now. Developments in this area have been significant. In contrast, energy management for wired networks (hubs, switches, routers) has not received much attention until recently. Gupta and Singh [13] were probably the first to bring these matters to light. These and other researchers have since then proposed several strategies to make routers and switches more energy aware. Most techniques consist in putting components to sleep or adapting link rate during periods of low traffic[16, 13, 8, 1].

Taking in consideration recent developments in this field and anticipating some of the future technological advances, in this paper we propose a scheme to reduce bandwidth and energy consumption for IPTV distribution networks. The IPTV services considered here target a TV viewing environment integrated with set-top boxes (STBs), providing cable TV-like experience. This service is offered on top of an IP network, instead of the hitherto predominant cable networks or free air broadcasts.

In order to guarantee the QoS required (i.e., to minimise channel change time), service providers opt to multicast all TV channels to all clients at all times (in other words, each node in the network pre-joins all multicast groups). This is provably inefficient. As soon as the number of channels surpasses the number of users at a certain access node, sending all channels becomes wasteful. Also, recent work [7, 18] has shown that channel popularity is highly skewed (following a zipf-like distribution). While a small number of channels is very popular, dozens or even hundreds of TV channels are very rarely watched. Service providers seem to recognise this problem and are already concerned with the efficiency of their IPTV networks[6].

By analysing our dataset - a large number of traces from a commercial nationwide IPTV service, scaling to 255 thousand users and 150 TV channels -, we have realised that indeed for the most popular channels there is always at least one viewer in any access node (a DSLAM in the network analysed), at any one time¹. However, the number of active channels in each DSLAM is rarely above 60. This can be observed in Figure 1, where the number of active channels in a node (DSLAM or regional router²) are presented, as a function of the number of users. It can be observed that the network is provably wasting resources by distributing all 150 TV channels everywhere. In the figure, the shaded area represents the bandwidth saving opportunities for an IPTV service provider. With the expected increase of the number of channels (AT&T, for instance, is already offering more than 700 TV channels[17]) this inefficiency will only worsen.

¹In this paper a channel that has at least one viewer in a node (DSLAM or router) is called an **active channel** in that particular node.

²In this paper we use the terms regional, metro, and edge network interchangeably.




Figure 1: Relationship between number of users and number of active channels

Considering the above, we propose to reduce these inefficiencies by not distributing all IPTV channels to everywhere at all times. Instead of pre-joining all multicast groups, we propose to pre-join only a selection of channels. We designed our scheme with two goals in mind: 1) it should be simple to implement, and 2) it should be dynamic. This second point is relevant, since users from different regions may have different preferences, and these preferences may change with time (the dynamic nature of channel popularity was analysed in [18]). We achieve these two design goals by each node (be it a core, edge, local router or a DSLAM) pre-joining only the active TV channels plus a small subset of the inactive ones. This scheme is very simple to implement and is dynamic, since the list of pre-joined channels varies with user activity.

We evaluate our scheme in two ways. Firstly, we analyse the bandwidth savings and contrast these with the number of requests affected. A request is considered affected when the user requests a channel that is not in the selection list (i.e., that particular node has not pre-joined that particular multicast tree). In such case, the user will experience a zapping (or channel switching) delay higher than usual. This is due to the fact that the join message to the multicast group will have to go some levels up to the next nearest branch of the multicast tree. In Section 4 we explain why this will not compromise user experience seriously. Secondly, we analyse the impact these bandwidth savings will have on energy consumption. For this purpose we create a power consumption model for two routers, based on real measurements[8], and evaluate it considering present and futuristic scenarios.

To evaluate our scheme we performed trace-driven simulations on the dataset described above. The results show that it is possible to significantly reduce the bandwidth used in the network, affecting only a very small number of switching requests. Today, these savings will have negligible impact on energy consumption. In the future, however, energy savings will start to become meaningful. We anticipate that the expected developments in network equipment will have difficulty to cope with the increase of very high quality IPTV bandwidth demands.

Our results are in line with the recent conclusion that the energy bottleneck will move to the core[21], and that energy consumption will be dominated by the routing function in the future. For this reason, IPTV service providers will need to reconsider their IPTV distribution networks.

2. DESCRIPTION OF AN IPTV NETWORK

Incumbent operator's IPTV networks are "walled gardens", well provisioned to guarantee the user experience required by TV viewers. Current networks use static IP multicast within a single network domain. The TV head-end injects live TV streams encoded as IP packets to the network core. These packets travel through the network to routers that are connected to the access nodes (DSLAMs if it is a DSL network, as in the network here analysed). These DSLAMs aggregate traffic from several users (usually in the order of the hundreds) connecting them to the high-speed backbone. The TV channels are distributed from the TV Head-end to the DSLAMs through bandwidth-provisioned, static, multicast trees. The network is dynamic exclusively between DSLAMs and customers, where multicast trees are extended or pruned based on the channel switching signals. The Set Top Box translates a channel switching event from a user remote control into a pair of Internet Group Management Protocol (IGMP) messages: one to inform of the user's departure from the multicast group of a TV channel, and the other to join another multicast group of the new channel.

3. DESCRIPTION OF THE TRACES

The evaluation of the scheme presented in this paper is made by trace-driven simulations. The dataset used for this purpose is from a commercial, nationwide, IPTV service. The logs were collected over a six month period. The traces scale to 255 thousand users, 623 DSLAMs, 150 TV channels, and 11 regions.

The trace includes all switching events during this 6-month period, and each event comprises the following information: timestamp in units of seconds; DSLAM ID; Set-top-box (STB) ID; IP address of the multicast group (TV channel); Type of switching event: join or leave channel. Note that all data was anonymised for the purpose of this study. No information that could directly or indirectly identify individual subscribers was included.

4. SELECTIVE PRE-JOINING

Currently, IPTV operators distribute all IPTV channels to everywhere at all times, in order to minimise channel switching (zapping) delay. This means that all DSLAMs are leafs of the multicast tree of every TV channel. Therefore, multicast nodes in the network pre-join all TV channels. Our proposal is for each node to pre-join only a subset of the full selection of TV channels. Namely, each node will join a) the channels for which there are viewers (the active channels) plus b) a small subset of inactive ones. We call the number of inactive channels that are pre-joined in the node the *room size*.

Our scheme requires two simple data structures: one to maintain information on the pre-joined channels, and another to record if each of these is active or inactive. When a TV channel A becomes inactive, one of two things can happen: 1) if the number of inactive channels is smaller or equal to the room size, keep channel A in the list of pre-joined channels; 2) if the number of inactive channels is bigger than the room size, then we randomly remove one of the inactive channels, by sending a leave message towards the source.

If the user requests a channel that is not active, the request has to go up to the nearest branch of the multicast tree. It

will therefore experience a larger than usual delay, due to an increase in network delay. This network delay is a small contributor to the overall zapping delay in IPTV: buffering and stream synchronisation are the main culprits[20]. Also, if we satisfy one of the objectives of our scheme - to affect a very small number of channel requests - the number of signalling messages transported in the network will not be significant, and hence overall network delay should not be affected.

A relevant point to emphasise is that this scheme "automatically" pre-joins the channels that have an higher probability of being watched in the future. Popular channels are the ones people watch more, so these are always pre-joined. Since the pre-joining list maintains the active channels at any given moment, the selection changes dynamically with user demand.

4.1 Evaluation

To evaluate our scheme we have performed trace-driven simulations on the large dataset presented before. All results presented in this section arise from the analysis of the full data set. In Figure 2 we present a graph with the bandwidth savings achieved with the several schemes tested. The figure presents the median, 10th- and 90th-percentile over the whole dataset. We analysed two types of nodes: local (access/DSLAM level) and regional (edge routers). At the access network level we tested our scheme for five different values of room size. With a room size of 20, bandwidth can be reduced by 50%. At the regional level we pre-joined only the active channels, so the room size was zero. The average bandwidth savings were equal to 33%. We compare our schemes with the one currently used by network providers ("All channels").




Figure 2: Bandwidth savings

In order to analyse the effect these schemes will have on the QoS, we inspected the percentage of requests to channels not pre-joined in the node. This is shown in Figure 3. For a room size of 20 (in the figure, r represents room size), the percentage of requests affected at the DSLAM level is less than 2%. At the regional level this figure is almost negligible. Therefore, we believe using a room size of 20 at the access level offers interesting results. In regional routers, pre-joining active channels only seems a good option.

The bandwidth savings mentioned before correspond to a reduction in the number of TV channels pre-joined. By pre-joining less channels, the nodes will have to process less bits, and the links immediately above in the network hierarchy will transport less bits. How significant these savings are, both in resource and energy terms, is intrinsically dependent on the bitrate at which channels need to be distributed. To




Figure 3: Percentage of requests affected

Table 1: Description of scenarios

Scenario	Media format	Bit rate	TV channels	Bandwidth savings
Today	SDTV	4 Mbps	150	0.3 Gbps
In 2-4 years	HDTV	20 Mbps	700	7 Gbps
In 10-15 years	4K	200 Mbps	3000	300 Gbps

clearly understand the significance of the savings achieved by our scheme, we analyse three scenarios, characterised in Table 1.

In the network just analysed 150 TV channels are distributed in Standard Definition format (SDTV). A bandwidth saving of 50% means reducing load by 300 Mbps. This is not very significant. Currently, however, most IPTV networks already offer more channels (AT&T offers 700[17]) with many in high definition (HDTV). Assuming such scenario we are now talking of bandwidth reductions of around 7 Gbps. Looking a bit further into the future, we can anticipate even higher quality streams (digital cinema standard 4K, for instance, or UHDTV[12]). In addition, according to a recent press release by the European Commission[11], over one thousand channels have been established in the UK alone until 2009, so assuming 3000 TV channels are offered in 10 to 15 years seems reasonable. In this case, bandwidth savings would be very significant, with a magnitude of several hundreds Gbps.

The scale of our dataset offers confidence in the results obtained. Furthermore, we believe this scheme would achieve similar gains in other IPTV networks. This is because the basis of our scheme is the observation that only a subset of TV channels are being watched at any one time, and this is rooted in a specific fact: channel popularity is highly skewed (zipf-like), meaning that some channels are very popular, while others are very rarely watched. This fact was observed by different research groups using different datasets: one group [7] analysed the same dataset used in this study (150 TV channels, 250k users), and another group [18] analysed a different dataset with a different scale (700 TV channels, 1 million users). Both reached the same conclusion in terms of channel popularity, so we are confident that our analysis can be generalised.

5. IMPACT ON ENERGY CONSUMPTION

Saving bandwidth may be, *per se*, an important objective. In this paper, however, we additionally proposed to analyse the impact of our scheme on energy consumption. In principle, bandwidth savings should result in energy savings. Less bits need to be transported in the links, and less bits need to be processed in the routers. Reducing load in the network offers more opportunities to put some equipment to sleep or to adapt line rates, in order to save energy. In this

section we will try to understand if these bandwidth savings are translated in relevant energy savings.

Unfortunately, there is a lack of power measurement studies from live networks. Besides this, there is no good understanding of how the energy consumed by network equipment varies under different loads and router configuration settings. For this reason, we will devise a power consumption model based on one study that offers real measurements of existing network equipment [8].

5.1 Power consumption model

Energy consumed by a router depends on five variables: 1) power to maintain the chassis on: a fix amount independent of load; 2) number of active linecards; 3) number of active ports in each linecard; 4) port capacity; and 5) port utilisation.

Considering this, in our model we assume that the routers: 1) will be always on; 2) support a mechanism to turn off line cards; 3) do not support a mechanism capable of turning off specific ports in an active linecard; 4) cannot change port capacity; and 5) the energy usage per port is independent of utilisation. We also assume that any pair of routers will maintain multi-bonded channels to inter-communicate. This is commonly the case, with the links between pairs of routers transporting multiple parallel channels, either through multiplexing or by using multiple optical fibres. Finally, we assume that the routers will use each of these parallel channels to its full capacity before deciding to use a new free channel.

This is a conservative model because assumes only that linecards can be turned off automatically when not needed. We expect in the near future routers to not only being capable of doing this, but also of automatically turning off specific ports from an active linecard and generically have a more energy proportional behaviour (the IEEE, for instance, already has a working group devoted to these matters[14]).

Based on these assumptions, we consider the two routers analysed in [8]. We assume the Cisco 7507 equipped with six 1GE linecards, and the Cisco GSR 12008 equipped with eight 4x1GE. The model developed for our analysis is shown in Figure 4, with the inputs defined in Table 2.




Figure 4: Model of power profile

As one line card is turned on, the power consumption increases by Δ_p , hence the discrete-steps increase. The energy per bit figures included (from [8, 21, 9]) allow us to extrapolate the power consumption profiles of the more advanced Cisco CRS-1 and Cisco CRS-3 routers in the next section.

5.2 Evaluation of different scenarios


We will now analyse the impact the bandwidth savings

Table 2: Model inputs

Router	Δ_p (W)	Δ_l (Gbps)	Energy per bit
Cisco 7507	30 [8]	1	65nJ/bit [8]
Cisco GSR 12008	92 [8]	4	36nJ/bit [8]
Cisco CRS-1	N/A	N/A	10nJ/bit [21]
Cisco CRS-3	N/A	N/A	3nJ/bit [9]

reported in Table 1 would have in energy consumption. We consider again those three scenarios. For all scenarios, we assume the network is composed of 250 edge routers and 50 core routers. Intel network has approximately 300 routers and was used as reference in earlier energy-related work[16], so we assume this to be also a representative figure for a nationwide IPTV provider. In the first scenario, “today”, we assume the edge routers to be Cisco 7507, while the core uses Cisco GSR 12008. In the futuristic scenarios we need more powerful routers, so we assume for the second scenario Cisco GSR 12008 in the edge and Cisco CRS-1 in the core, and for the medium-long term scenario CRS-1 in the edge and CRS-3 in the core. Since we do not have details from real measurement on these routers, we will extrapolate their Δ_p and Δ_l values. We will assume the same values as for the GSR 12008, but considering also the increased power efficiency of these routers, represented by the energy per bit figures. Specifically, we will assume the CRS-1 to be three times as efficient as the GSR 12008 (based on [21]), and the CRS-3 to be three times as efficient as the CRS-1 ([10]).

The energy and cost savings for all three scenarios are presented in Figure 5. We present the results in terms of reduced power consumption (in kWh per year) and reduced cost (in € per year, considering 0.1€ per kWh).

**Figure 5: Energy and cost savings**

Today, the bandwidth savings achieved with our scheme will have negligible impact on energy consumption, at least in the core network. Nevertheless, it is important to note that we have not considered the access network. An ADSL2+ DSLAM, for instance, has an energy per bit figure two orders of magnitude higher than that of the Cisco 7507 router[2]. It is therefore quite probable that reducing the number of bits that need to be processed at this level offers significant energy savings. Unfortunately, we do not have data to create a power profile model for this equipment, so this remains a supposition.

In the future, however, the energy savings will start to become meaningful, even considering our conservative power profile model. This is grounded in the observation that the energy per bit figure in network equipment has been reduced by around one order of magnitude per decade in the past and we expect it to follow a similar trend in the future. By considering an increase of IPTV network bandwidth demands

by three orders of magnitude, the energy footprint caused by this type of service may become significant.

6. CONCLUSIONS AND FUTURE WORK

Delivering TV streams using an IP network means consuming a significant amount of the network’s resources. As the number of TV channels increases and the quality of the streams improves (with the resulting increase of the bandwidth requirements), a network provider will have to think carefully if it can handle that burden. Fortunately, the majority of users tend to enjoy the same TV channels: 90% of all TV viewing is restricted to a small selection of channels[7, 18]. We argue, therefore, that an IPTV provider can use this fact for its own benefit. Instead of broadcasting all channels at all times, it can judiciously choose which TV channels to distribute where at any one time.

In this paper, we propose each network node to pre-join only a selected list of these channels: the active TV channels plus a small subset of the inactive ones. We evaluated this scheme by performing trace-driven simulations on a large dataset from a nationwide IPTV network. We contrasted the bandwidth savings achieved with the number of requests affected, and concluded that it is possible to reduce bandwidth significantly without compromising service quality.

A power consumption model was developed to assess the impact these bandwidth savings will have on energy consumption. The main conclusions were that nowadays energy savings in the core are not significant, in the not improbable scenario of IPTV bandwidth demands increasing faster than electronic routers energy-efficiency developments, network inefficiencies will become an issue deserving consideration.

As future work, we will continue investigating techniques to improve the efficiency of IPTV networks. We intend to mature this work, by refining the power consumption model and by performing detailed sensitivity analysis. We will take into consideration the expected evolution of routers in the next decade. We are also planning to consider optical bypass as a strategy to make the growth of IPTV network power consumption more manageable in the future. Using this technique for long-lived flows (popular channels are a good target), we reduce router work (less bits to process) and make good use of the low power consumption properties of optical cross connects (switching energy of an all optical circuit switched cross connect is lower by two orders of magnitude when compared to its electronic equivalent [2]).

7. REFERENCES

- [1] Y. Agarwal, S. Hodges, R. Chandra, J. Scott, P. Bahl, and R. Gupta. Somniloquy: Augmenting network interfaces to reduce PC energy usage. In *Proc. NSDI*, 2009.
- [2] J. Baliga, R. Ayre, K. Hinton, W.V. Sorin, and R.S. Tucker. Energy consumption in optical IP networks. *JLT*, 27/13:2391–2403, July 2009.
- [3] J. Baliga, R. Ayre, K. Hinton, and R.S. Tucker. Photonic switching and the energy bottleneck. In *Photonics in Switching*, 2007.
- [4] J. Baliga, K. Hinton, and R.S. Tucker. Energy consumption of the Internet. In *Proc. COIN-ACOFT*, 2007.
- [5] L.A. Barroso and U. Holzle. The case for

- energy-proportional computing. *IEEE Computer*, 40/12:33–37, December 2007.
- [6] J. Caja. Optimization of IPTV multicast traffic transport over next generation metro networks. In *Proc. NETWORKS*, 2006.
 - [7] M. Cha, P. Rodriguez, J. Crowcroft, S. Moon, and X. Amatriain. Watching television over an IP network. In *Proc. ACM IMC*, 2008.
 - [8] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsang, and S. Wright. Power awareness in network design and routing. In *Proc. IEEE INFOCOM*, 2008.
 - [9] Cisco. <http://www.cisco.com>.
 - [10] CISCO. Cisco visual networking index: Forecast and methodology 2008-2013, 2008.
 - [11] European Commission. Growth of the number of TV channels and multi-channel platforms in Europe continues despite the crisis, 2010.
 - [12] M. Ghanbari, D. Crawford, M. Fleury, E. Khan, J. Woods, H. Lu, and R. Ravazi. Future performance of video codecs, 2006.
 - [13] M. Gupta and S. Singh. Greening of the Internet. In *SIGCOMM*, 2003.
 - [14] IEEE. P802.3az energy efficient ethernet task force. <http://www.ieee802.org/3/az/>.
 - [15] C. E. Jones, K. M. Sivalingam, P. Agrawal, and J.-C. Chen. A survey of energy efficient network protocols for wireless networks. *Wireless Networks*, 7:343–358, 2001.
 - [16] S. Nadevschi, L. Popa, G. Iannaccone, S. Ratnasamy, and D. Wetherall. Reducing network energy consumption via sleeping and rate-adaptation. In *Proc. NSDI*, 2008.
 - [17] T. Qiu, Z. Ge, S. Lee, J. Wang, J. Xu, and Q. Zhao. Modeling user activities in a large IPTV system. In *Proc. ACM IMC*, 2009.
 - [18] T. Qiu, Z. Ge, S. Lee, J. Wang, Q. Zhao, and J. Xu. Modeling channel popularity dynamics in a large IPTV system. In *Proc. ACM SIGMETRICS*, 2009.
 - [19] V. Raghunathan, M. B. Srivastava, and R. K. Gupta. A survey of techniques for energy efficient on-chip communication. In *Proceedings of Design Automation Conference*, pages 900–905. ACM Press, 2003.
 - [20] P. Siebert, T.M.N.V. Caenegem, and M. Wagner. Analysis and improvements of zapping times in IPTV systems. *IEEE Trans. on Broadcast.*, 55/2:407–418, June 2009.
 - [21] R. S. Tucker, R. Parthiban, J. Baliga, K. Hinton, R. W. Ayre, and W.V. Sorin. Evolution of WDM optical IP networks: A cost and energy perspective. *JLT*, 27/3:243–252, Feb. 2009.