Blending photons with electrons to reduce the energy footprint of IPTV networks

Fernando M. V. Ramos*, Jon Crowcroft†, Ian H. White†
LaSIGE, Faculdade de Ciências, Universidade de Lisboa, Portugal* University of Cambridge, UK†

***Abstract*—The rapid growth of IPTV services and the resulting increase in traffic volumes is raising concerns over energy consumption. In this paper we propose to save energy by shifting particular IPTV traffic from power-hungry electronic routing to greener optical switching. The traffic profile of IPTV results in such a hybrid switching approach to allow *both* energy and bandwidth efficiencies. To achieve this goal we designed a novel protocol that allows the use of optical bypass in IPTV networks. By means of a trace-driven analysis of a large dataset we demonstrate the energy efficiencies obtained to be substantial, reaching power savings of over 40% under normal load conditions. This result represents a four-fold increase in energy efficiency when compared with recent proposals.**

I. Introduction

The past few years have witnessed the rapid roll-out of IPTV services. IPTV has been launched by major service providers worldwide [1] and its popularity is on the rise. According to the most recent report from Digital TV Research [2], in the period from 2011 to 2013 the number of IPTV customers has increased from 36 to over 90 million, and is expected to reach 191 million by the end of 2020.

IPTV is a resource intensive service with stringent quality of service requirements. 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 (4K, UHDTV). Besides the resulting increase in bandwidth requirements, the number of TV channels offered is also expected to grow. Current IPTV operators already offer near one-thousand TV channels [3] to its customers. But recent trends anticipate the likely growth of the number of TV channels in the near future. Narrowcasting services – broadcasting to a very small audience [4] –, for example, are growing in importance.

The increase in bandwidth required to support an increasing number of users and of TV content has led to concerns about the energy consumed by the infrastructure. Various studies have highlighted the effects of GreenHouse Gases (GHG) emissions and their consequences on climate change [5] and on the economy. ICT represents an important source of energy consumption and of GHG emissions, with 37% of the total ICT emissions due to the network infrastructure [6].

In this paper we focus our attention in the energy consumption of traditional (push-based) linear TV over IP networks. In this type of system TV programs are broadcast on the different channels according to a known schedule (content is therefore *pushed* to users). This type of service is fundamentally different from pull-based approaches such as time-shifted TV[1] or VoD. Linear TV broadcasts do not have the natural scaling properties of time-shifted TV [7], increasing the challenge of improving the efficiency of these networks. For instance, the intrinsic nature of linear TV services precludes the use of energy optimisation approaches based on caching techniques, as recently-proposed for time-shifted TV [7], [8]. Alternative techniques are therefore needed.

Today, IPTV systems[2] are inefficient. In current deployments *all* TV channels are distributed to *all* local routers, despite particular channels having no viewers at particular time periods. Considering this inefficiency problem, we have proposed a scheme [9] – *selective pre-joining* – where only a selection of TV channels is distributed in the network, instead of all. We have shown, based on real data, that by using this scheme it is possible to save bandwidth and energy while affecting a very small number of channel switching requests. For instance, we have demonstrated that if core routers pre-joined only the channels that have viewers the effect in the quality of service would be residual (less than 0.1% channel requests would be affected) while achieving important efficiency gains.

In this paper we propose to go further in terms of energy-efficiency by following a different approach. The technique we propose is based on the introduction of optical switching in the network. The rationale is the fact that optical switching techniques are more energy-efficient than their electronic counterpart [10]. In particular, we assess the opportunities for performing *optical bypass* in IPTV networks, and propose a novel protocol for this purpose. With optical bypass, traffic not destined for a given network node is not processed electronically by that node. This traffic is all-optically switched.

By avoiding electronic processing and performing optical switching instead, energy savings are to be expected. We demonstrate in this paper this hypothesis to be true. For this purpose, we evaluated our protocol by means of trace-driven analyses using a dataset from an operational IPTV service provider. The dataset scales up to 150 TV channels, six months, and 255 thousand users. We demonstrate that by using the proposed scheme IPTV service providers can significantly reduce the energy consumption of their networks. For instance, for normal traffic load conditions our proposal presents a four-fold increase in energy-efficiency when compared with *selective pre-joining* [9].

[1] On-demand access to previously broadcast TV content, also known as catch-up TV.

[2] We will henceforth use IPTV to refer to linear TV services distributed over an IP network.

ISBN 978-3-901882-83-8 © 2016 IFIP

II. Background

A. IPTV networks

A traditional "walled garden" IPTV network can be split logically into three main domains – the access network, the metropolitan network, and the IP network. The IP network usually has a two-level, hierarchical structure [11]: the regional network (sometimes called edge) and the core (Figure 1).

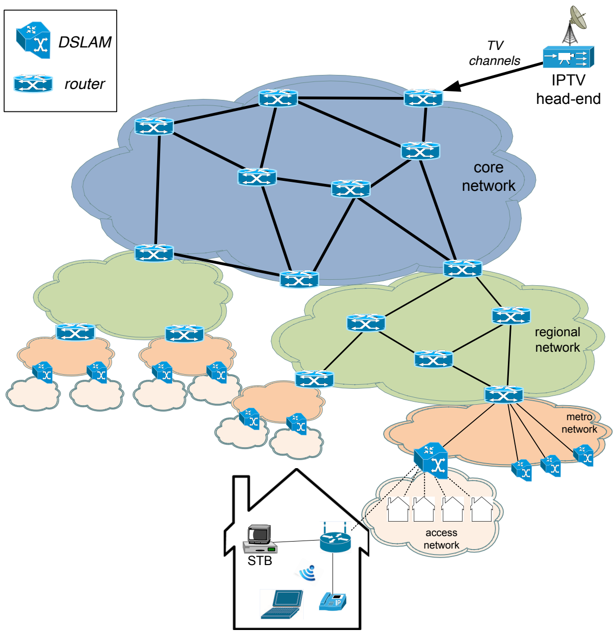

Fig. 1. IPTV network

In an IPTV system, live TV streams are encoded in a series of IP packets and delivered through the network to the residential broadband access network. The core network comprises a small number of large routers in major population centres. The core routers are often highly meshed, with high-capacity WDM fibre links interconnecting them. In the regional network, routers are normally lower-end routers with high port density, where IP customers get attached to the network. These routers aggregate the customer traffic and forward it toward the core routers. The metro network serves as the interface between the regional network and the access network. The access network connects each home to one of the edge switches in the provider's network. There is a wide variety of access technologies: from ADSL (Asymmetric Digital Subscriber Line) to fibre-based solutions (FTTx). Inside the household, a residential gateway connects to a modem and one or more Set Top Boxes (STBs). Finally, each STB connects to a TV.

The TV channels are distributed from the TV head-end to edge nodes (DSLAMs in Figure 1) through bandwidth-provisioned multicast trees. Current networks use *static* IP multicast within a single network domain. By static multicast we mean all receivers are known beforehand, and no new group members are allowed to join – we have a *static set of receivers* for all TV content. Again referring to Figure 1, this means *all* DSLAMs join *all* multicast groups (thus receive content from *all* TV channels). This occurs despite the fact that particular channels have no viewers at particular time periods [9].

B. Core optical IP networks

An optical IP network can be seen as being made up of two layers, the IP layer and the optical layer. This is shown in Figure 2. In the first generation of optical networks, all the lightpaths incident to a node had to be terminated, i.e., all the data carried by the lightpaths would be processed and forwarded by IP routers. This is represented in the figure by lightpath 1. This wavelength is OEO (Optical-Electrical-Optical) converted at each node. In contrast, the new generation of optical networks includes elements such as the Optical Cross Connect (OXC) which allow some lightpaths to bypass the node. This approach allows IP traffic whose destination is not the intermediate node to directly bypass the intermediate router via a cut-through lightpath. This is represented by lightpath 2. This wavelength bypasses all nodes.

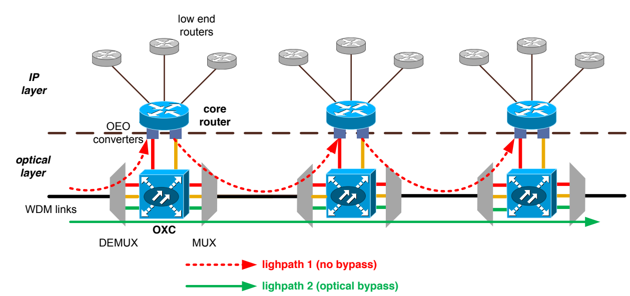

Fig. 2. Optical network employing optical bypass techniques

Several researchers have pointed out that optical bypass technology is one important method to reduce the power consumption of IP networks [10], [12], [13]. This technique can save energy because it can reduce the total number of active IP router ports, and these play a major role in the total energy consumption of an optical IP network [13]. Shifting traffic from power-hungry routers to low-power optical switches by means of optical bypass is therefore an effective technique to save energy in optical networks.

III. Optical bypass of popular TV channels

With the goal of reducing energy consumption, in this paper we consider the introduction of energy-friendly optical switching techniques in the core of optical IPTV networks. In particular, we propose the introduction of optical bypass. With optical bypass traffic not destined for a given IP router is placed onto a wavelength that is not processed by that router. Instead, this traffic is all-optically switched. Due to the circuit-switching nature of optical networks, however, only *long-lived flows* can be considered realistic targets for optical bypass.

Conveniently, some IPTV traffic is in this category. Some TV channels are very popular, having viewers everywhere in the network, at any particular time. Optically switching such long-lived flows can therefore be advantageous energy-wise. Other less popular and niche channels have periods without any viewers in particular locations, so it is wasteful to distribute them continuously everywhere. The dynamic nature of electronic packet-switching nodes is therefore ideal to switch this type of traffic. This guarantees the network is bandwidth efficient, by allowing these TV channels to be quickly removed from or added to the network as needed.

A. Protocol for optical bypass in IPTV

Considering the above, we propose a protocol to be used in the core of IPTV distribution networks, blending electronic routing with all-optical switching. We assume the network core to be composed of hybrid nodes, each including a multicast-capable WDM optical cross connect (OXC) and a multicast-enabled IP router, as illustrated in Figure 2. The inclusion of the OXC between the input ports and the router allows optical bypass to be performed. We further assume these nodes to be GMPLS-capable. A unified control plane such as GMPLS allows the integration of optical circuit-switching techniques with electronic packet-switching. The main idea of the scheme is for popular TV channels to be all-optically switched (switched at the optical layer), while the rest are electronically routed (switched at the IP layer). The network distributes the two different groups of channels in two (disjoint) sets of wavelengths. The wavelengths from one set optically bypass the nodes, whereas the other wavelengths are sent to the routers for processing. We restrict the use of the proposed scheme to the optical network core, as this is the only location where it is realistic to assume the presence of OXC equipment in the medium-term.

Algorithm 1 Processing at the IPTV source

1: **while** true **do**
2:   sleep($\Delta\tau$)
3:   send_to_core-reg_nodes(ACTIVE_CHANNELS_REQ)
  {Wait until all requests are received...}
4:   $CPop \leftarrow$ ALL_TV_CHANNELS
5:   $CNonPop \leftarrow \emptyset$
6:   **for** $i = 1$ to NUMBER_OF_NODES **do**
7:     $CPop \leftarrow CPop \cap ActiveCh[i]$
8:   **end for**
  {$CPop$ now includes all popular TV channels}
9:   **for** $i = 1$ to NUMBER_OF_NODES **do**
10:     $CNonPop \leftarrow CNonPop \cup (ActiveCh[i] \notin CPop)$
11:   **end for**
  {$CNonPop$ now includes the other TV channels with viewers}
12:   $\lambda_o \leftarrow$ [Wavelengths filled with $CPop$ channels]
13:   $\lambda_e \leftarrow$ [Wavelengths filled with $CNonPop$ channels]
14:   send_to_all_nodes(SWITCH_CHANGE_REQ, λ_o, λ_e)
15: **end while**

Algorithm 2 Processing at each core-regional node

1: **while** true **do**
2:   MESSAGE = msg_rcv_from_source()
3:   **if** MESSAGE == ACTIVE_CHANNELS_REQ **then**
4:     $ActiveCh \leftarrow$ get($McastFwdTable$)
5:     send_to_source($ActiveCh$)
6:   **end if**
7: **end while**

The protocol for optical bypass in IPTV networks consists of three algorithms. Algorithm 1 runs at the IPTV source, Algorithm 2 runs at core-regional nodes (Figure 1), and Algorithm 3 runs at the core nodes (including core-regional ones). The details of the proposed protocol follows.

Step 1. After a specified time interval, $\Delta\tau$, the source transmits a message requesting all hybrid core-regional nodes to submit their active channels (algorithm 1, lines 2-3). An *active channel* is a channel for which there is at least one viewer. This message sent by the source serves as a trigger for all core-regional routers to send this information back to the source as soon as possible. Considering that all nodes are GMPLS-capable, this information can be sent as a TE Notify message. RSVP-TE Notify messages were added to RSVP-TE[3] to provide general event notification to nonadjacent nodes.

Algorithm 3 Processing at each core node

1: **while** true **do**
2:   MESSAGE = msg_rcv_from_source()
3:   **if** MESSAGE == SWITCH_CHANGE_REQ **then**
4:     **for all** $\lambda \in \lambda_o$ **do**
5:       switch_optically(λ)
6:     **end for**
    {Wavelengths in the set λ_o are optically bypassed}
7:     **for all** $\lambda \in \lambda_e$ **do**
8:       route_electronically(λ)
9:     **end for**
    {Wavelengths in the set λ_e are sent to the router}
10:   **end if**
11: **end while**

Step 2. Each regional-core node then sends information on its *active channels* to the IPTV source. As the active channels are those being distributed by the regional-core router to its region, the multicast forwarding table of this router contains a line with their multicast group addresses and the interfaces used to forward packets to. The information requested can thus be easily retrieved and sent back to the source (algorithm 2, lines 3-6). Again, an RSVP-TE Notify message can be used.

Step 3. Once the source receives these sets from all routers, it checks which TV channels should be optically switched (the popular ones), and which should be electronically routed (the remainder channels with viewers). The popular channels are those which have viewers everywhere. Their multicast group addresses are present in the multicast forwarding tables of every core-regional router. The intersection of all sets received by the source thus results in a new set with the list of popular channels[4] (algorithm 1, lines 6-8). The union of the active channels of each set which are not popular results in a set with the non-popular TV channels (algorithm 1, lines 9-11).

Step 4. The TV channels are distributed, from the source, in two distinct sets of wavelengths: λ_o and λ_e. The popular channels are distributed using N different wavelengths: $\lambda_o = N \times \lambda$. The others are sent in a disjoint set of M different wavelengths: $\lambda_e = M \times \lambda$. The number of wavelengths in each set depends on the number of TV channels and its bit rate, and on the capacity of each wavelength. The IPTV source decides the composition of each set of wavelengths and informs all core nodes of its decision (algorithm 1, lines 12-14). This information can be sent in the form of an RSVP-TE PATH message. This is one of the messages used to allocate resources in the network. In multicast scenarios, only one PATH message

[3]As its name implies, the Resource Reservation Protocol - Traffic Engineering (RSVP-TE) is an extension of the resource reservation protocol (RSVP) for traffic engineering, and is used as part of the GMPLS control plane for this purpose.

[4]We are abusing the term "popular" in this paper. If one TV channel has a single viewer in each region then it is included in the popular set. We use this term to ease the understanding of the scheme.

ISBN 978-3-901882-84-5, IFIP Networking 2016 © 2016 IFIP

needs to be sent to multiple receivers, thus conserving network bandwidth.

Step 5. Each core node then sets up its switching state to optically switch the λ_o group (these wavelengths will therefore optically bypass the routers), and electronically route the λ_e group (algorithm 3, lines 3-10).

IV. Methodology and Dataset

The research community working on IPTV systems has relied upon hypothetical user models which are sometimes different from reality and can lead to incorrect estimation of system performance. Constant-rate Poisson models are generally used as workload model for these systems. Examples include [14], [15], among several others. Unfortunately, these models does not capture IPTV user behaviour well. For instance, users switch channels more frequently than these simple models predict. This fact was proved by Qiu *et al.* [3]. The authors made a comprehensive analysis of real data from an operational nationwide IPTV system (AT&T) where they show that the simple mathematical models generally used are not good. Faced with this concern, the solution we propose in this paper is evaluated by means of trace-driven analysis of a large dataset from an IPTV provider.

The analysis of our dataset led us to the same conclusions as in [3]. In Figure 3 we exemplify one of the problems of using a simple Poisson distribution as a mathematical model to represent the behaviour of IPTV users. The figure presents the Cumulative Distribution Function of the number of channel switches during one-minute periods (a *zapping period*, according to [1]). The analysis was done on the entire dataset we describe in Section IV-A. In the figure we compare the empirical data with a Poisson distribution with parameter λ equal to 1.948 (the one that fits better the empirical data).

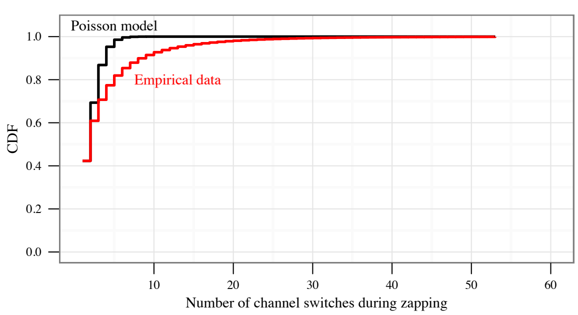

Fig. 3. Number of channel switches in zapping mode

As can be seen, the Poisson model is conservative in terms of the number of channel switches a user performs during zapping periods. For example, the probability of a user making five channel switches or more in a one-minute period is negligible when using the Poisson distribution. But in fact by observing the empirical data one can conclude that there is a 20% probability of a user switching channels five times or more during a zapping period.

A. Dataset

We obtained a collection of IPTV channel switching logs from an IPTV service offered by an operational backbone provider. This is a commercial, nationwide service, offering 150 TV channels over an IP network. The access links use ADSL technology and the network is composed of 680 DSLAMs distributed along 11 regions. To give an idea of the scale of the dataset, the 700GB trace spans six months and records the IGMP messages on the channel changes of around 255 thousand users. The number of daily channel switchings clocks 13 million on average. Table I summarises these statistics.

TABLE I. Dataset statistics

Trace duration	6 months
Number of users	255 thousand
Number of DSLAMs	680
Average number of daily channel switching events	13 million
Size of the dataset	700 GB

B. Validation of the dataset

To assure the representativeness of our dataset and the evaluation that ensues, we compared the results from Qiu et al [3] with the results from the analysis of our dataset (the two datasets are from different IPTV services offered in different countries). For validation we analysed the number of online users during the course of representative weeks, and found the same very strong diurnal patterns as Qiu *et al.* [3]. We also examined the long term distribution of channel popularity and found the same high skewness of popularity, which can be modelled using Zipf-like distributions. Finally, we also conclude that IPTV users switch channels very similarly in both studies. By analysing the entire dataset we observed that 55% of all channel switching was linear, up or down to the next or previous TV channel (in constrast to more targetted switching). Qiu *et al.* [3] reported 56% in the AT&T dataset. The full detail of this validation is out of scope of this paper. We leave an in-depth analysis as future work.

V. Evaluation

For the reasons explained before, the scheme we propose is evaluated by means of a trace-driven analysis. The IPTV trace detailed in Section IV-A is used as input to the analysis performed. All results we present next arise from the analysis of the entire data set (6 months, 255 thousand users). The evaluation is threefold. First, we investigate the scalability of the protocol. Second, we analyse the opportunities for optical bypass when running the proposed protocol in the network under study. Finally, we analyse the impact the use of this protocol has in power consumption of the IPTV network.

A. Scalability

For a network protocol to be scalable it is important that it does not impose a significant processing overhead to the network nodes and that it does not add a great amount of signalling traffic to the network. By guaranteeing a relatively long update interval for the control information (the $\Delta\tau$ variable in the proposed protocol) it is possible to guarantee a low overhead to the nodes and to the network as a whole. On the other hand, to assure the best performance it is important that the network state[5] is consistent with network usage (in this

[5] In this context, the network state consists of the wavelength switching configuration at each node, and the set of TV channels transported in each wavelength group, λ_o and λ_e.

ISBN 978-3-901882-84-5, IFIP Networking 2016 © 2016 IFIP

particular case, it should reflect channel popularity). Having a short update interval marries with this objective.

It is known that channel popularity is relatively stable over short time frames, and that it becomes more dynamic when longer time frames are considered [16]. Regular updates may therefore not be needed. We analyse the henceforth called *TV channel churn rate* in the 11 core-regional nodes of this network to attest this. We compare the active TV channels at time τ with the active channels at time $\tau + \Delta\tau$, for different values of $\Delta\tau$. The number of channels that are different between the two sets in two consecutive periods is the TV channel churn rate. The results are shown in Figure 4, for each region, and for five values of $\Delta\tau$. The median of the channel churn rate over the whole period of the trace (6 months) is presented, with the lower and upper error bars representing the 5th- and 95th-percentile, respectively.

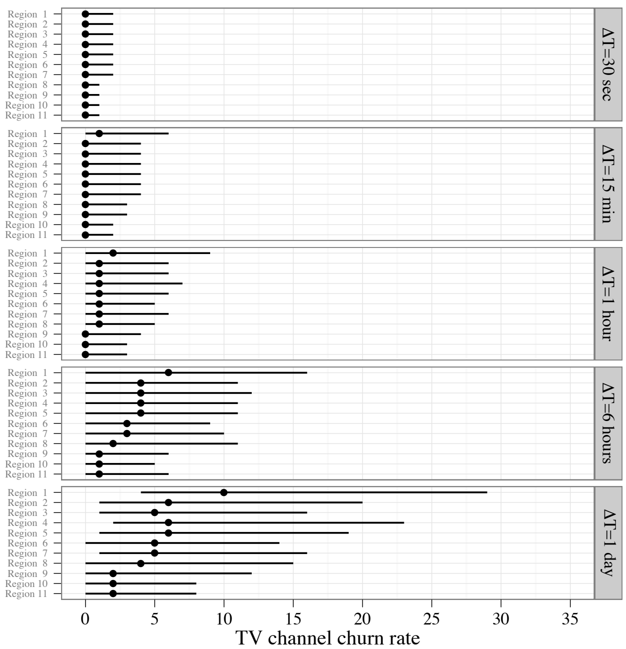

Fig. 4. TV channel churn rate for all eleven regions, for five values of the update interval

By analysing the results in Figure 4, we conclude that the churn rate is usually quite low, particularly for values of $\Delta\tau$ below 1 hour. A long update interval of 15 minutes, for instance, seems a good compromise. It does not represent a significant overhead to the network, while at the same time guarantees that the network state changes with channel popularity dynamics.

B. Opportunities for optical bypass

The protocol proposed divides the TV channels into three groups: the popular channels, the unpopular channels, and the channels without viewers. The channels from the former group optically bypass the routers. Those from the second group are sent for the router for electronic processing. Finally, those from the latter group are not distributed by the IPTV source (when a user switches to a channel without viewers a request is sent to the source for its quick distribution and the channel becomes part of the group of unpopular channels). To understand the opportunities for optical bypass in the core of the IPTV network, we need to quantify how many channels would be included in each group, at regular intervals. For this purpose, we retrieve the number of channels in each set (popular, unpopular, and channels with no viewers), for the entire trace. We consider for the analysis an update interval equal to 15 minutes, for the reasons explained above. This is the periodicity with which we retrieve the number of channels in each set. In Figure 5 we present the results obtained (median, 5th-, and 95th-percentile) from the analysis of the entire dataset.

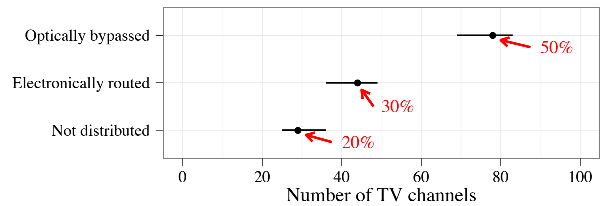

Fig. 5. Average number of TV channels that are optically bypassed, electronically routed and not distributed, respectively

We start the analysis from the bottom of the figure. On average, one fifth of the TV channels do not need to be distributed by the IPTV source. Recall that not distributing this traffic has a negligible impact on the service [9]. The remaining 80% TV channels are distributed to the network core. Around 50% of the TV channels can be optically bypassed. This means that, on average, at any one time, half of the channels have at least one viewer in each region. The number of channels requiring electronic processing can thus be reduced to around 30%.

C. Impact on energy consumption

After understanding that by using the proposed protocol there are clear opportunities to introduce optical bypass in IPTV networks we now analyse the impact this has on energy consumption. By employing this technique energy savings are expected for two reasons. First, some traffic flows (the popular TV channels) bypass some routers. This reduces the number of bits requiring electronic processing, thus avoiding energy-expensive OEO conversions, buffering, and forwarding table lookups. The work is shifted to optical switches, which are at least two orders of magnitude more energy efficient when compared to its electronic equivalent [10]. Second, as TV channels without viewers are not distributed, network load is reduced and even less bits require electronic processing in the routers.

1) Selective pre-joining in core optical networks: Before presenting results from our proposal, we return to the scheme proposed in [9], *selective pre-joining*, and use it as our baseline. However, we consider an optical network scenario, which was not considered by the original study. As in this work we consider a core optical network, we need to integrate the power consumption of the optical layer components into the model. This will allow a fair comparison with the proposal made by Ramos *et al.* [9]. In fact, this refined model will reinforce the effectiveness of *selective pre-joining*.

Several factors affect the power consumption of a core network node [17]. First, the base chassis power. This is the

TABLE II. LINECARD POWER PROFILE

Energy component and description	Estimate from [18].
Power consumed by unconnected linecard card (P_c)	6.936 W
Power consumed per connected Ethernet port (P_E)	1.102 W
Per-packet processing energy (E_p)	197.2 nJ
Per-byte energy (E_b)	3.4nJ

TABLE III. DESCRIPTION OF THE THREE SCENARIOS

Scenario	Media format	Bit rate	TV channels	**Bandwidth savings**
150SD	SDTV	4 Mbps	150	**0.3 Gbps**
700HD	HDTV	20 Mbps	700	**7 Gbps**
3kUHD	4K	200 Mbps	3000	**300 Gbps**

power to maintain the chassis on. It is a fix amount independent of load, including the power consumed by components such as fans, memory, etc. Second, the number of active linecards. A linecard is the electronic circuit that interfaces with the network. Third, the number of active ports in each linecard. Fourth, port capacity. This is the line rate forwarding capacity of individual ports. Fifth, port utilisation. This is the actual throughput flowing through a port, relative to its capacity. Sixth, power consumption of the transponders. In optical networks, associated with each wavelength (port) is a transponder (OEO converter), as was shown in Figure 2. The transponder interfaces the router to a fibre optic cable. Its main function is to perform the required OEO conversions. Considering this, the power consumption model is presented in Equation 1.

$$P = P_{ch} + \sum_{i=0}^{L} P_{l_i} + K_T P_T \qquad (1)$$

In this equation P_{ch} refers to the power consumption of the chassis. L is the number of linecards that are active, and P_{l_i} is the power consumption of linecard i. The power consumption of each linecard is calculated based on the model proposed by Sivaram *el al.* [18] for a NetFPGA card, and is presented as Equation 2. By using a high-precision hardware-based traffic generator and analyser, and a high-fidelity digital oscilloscope, the authors devised a series of experiments allowing them to quantify the per-packet processing energy and per-byte energy consumption of such linecard.

$$P_l = P_c + KP_E + N_I E_p + RE_b \qquad (2)$$

In this equation P_c is the constant baseline power consumption of the NetFPGA card (without any Ethernet ports connected); K is the number of Ethernet ports connected; P_E is the power consumed by each Ethernet port (without any traffic flowing); N_I is the input rate in packets per second (pps); E_p is the energy required to process each packet; R is the traffic rate in bytes per second (we assume the input rate is equal to the output rate); E_b is the total per-byte energy (this includes the energy required to receive, process and store a byte on the ingress Ethernet interface, and the energy required to store, process and transmit a byte on the egress Ethernet interface). The inputs to this model are presented in Table II, again based on the measurements reported in [18].

Finally, returning to Equation 1, K_T is the number of transponders (one per port) and P_T is the power per transponder. Every time a new port needs to be turned on, a new transponder is also activated. We assume the power consumption for each transponder to be 73 W, based on Alcatel-Lucent WaveStar OLS 1.6T ultra-long-haul systems [19]. This figure has been used in related work [13], [20].

For evaluation we consider the three scenarios presented in Table III: `150SD`, an IPTV service offering of 150 SDTV channels; `700HD`, 700 HDTV channels; and `3kUHD`, 3000 UHDTV channels. The first two scenarios represent current IPTV service offerings, whereas the latter is a futuristic scenario. For the first scenario, we assume a router with four linecards with 4x1Gbps Ethernet ports each. For the second scenario we scale up the node to sixteen linecards of the same type, for it to be able to handle the increased aggregate throughput. The capacity of each node is now assumed to be equal to 64Gbps. The capacity of the nodes of the third scenario has to scale up to the Tbps range. We assume fourteen 4x40Gbps linecards for an aggregate capacity of 2.2Tbps. This is a different type of linecard from the one measured by Sivaram *et al.* [18]. We therefore assume a 4x40Gbps linecard presents the same power profile as fourty 4x1Gbps.

The results we present first illustrate the relative power savings of using the *selective pre-joining* scheme as a factor of the baseline traffic load, according to equation 3. The baseline traffic load is the load of a node that does not use the scheme. This load obviously includes IPTV traffic.

$$\frac{P\,(baseline) - P\,(selective\_joining)}{P\,(baseline)} * 100 \qquad (3)$$

In Equation 3, $P\,(baseline)$ is the power consumption at baseline traffic load, whereas $P\,(selective\_joining)$ is the power consumption when using the proposed scheme (a lower value due to the decrease in IPTV traffic). In the figures we present results for baseline load values varying from 25% to 75%.

In accordance to the results presented in the previous section (Figure 5), we assume that only 20% of the channels are not distributed to the core. The results we present in Figure 6 thus correspond to a reduction of IPTV traffic in the network core to 80%. In this figure the lines labeled *(IP only)* are the results without considering optical components in the power consumption model (as per [9]), whereas *(IP + opt)* are the results using our augmented model, considering the optical transponders.

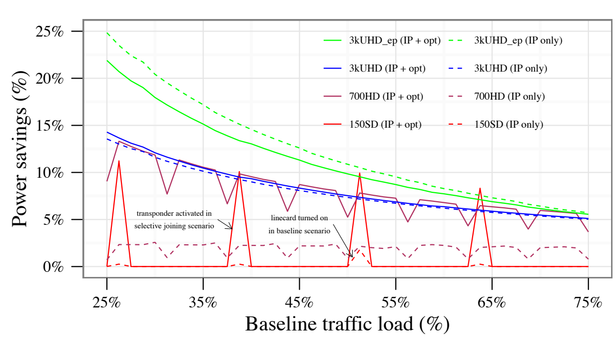

Fig. 6. Power savings of using the selective pre-joining scheme proposed by Ramos et al [9] considering an optical IP network (core network)

ISBN 978-3-901882-84-5, IFIP Networking 2016 © 2016 IFIP

As can be seen, by considering the optical ports the results change significantly. The main reason is the fact that the transponders are power-hungry equipment. This results in an increased advantage in using the *selective pre-joining* scheme in some scenarios, as reducing traffic load decreases the number of active transponders. It is particularly relevant to mention scenario `700HD`, which is typical in current networks (this scenario is based on AT&T's IPTV service offering [3]). The use of the scheme proposed in [9] increases the power savings to around 10% under normal traffic loads [11] when we consider optical components.

One aspect that deserves explanation is the lines in the plots not being completely smooth (the little "steps"). This is particularly evident in the first two scenarios, `150SD` and `700HD`. The reason is that the *x*-axis represents the baseline traffic load in the node (without *selective pre-joining*), while the power savings arise from the new traffic load (with *selective pre-joining*) being lower. The power saving peaks that appear in the graph represent transition points, when a particular event that increases significantly the energy consumption occurs: when an additional linecard needs to be turned on or the activation of another transponder (as this component consumes more power than a linecard, the peaks are more pronounced). For instance, in the `150SD` *(IP only)* scenario there is a peak precisely in the middle of the plot. This is because a 50% load in that scenario represents a data rate equal to 4Gbps. At this point, the network node has to turn on a new linecard (recall that we are assuming 4x1Gbps linecards). With *selective pre-joining* the network load would be lower than the baseline traffic load, a bit under 4Gbps. So the linecard does not need to be turned on yet. While the traffic load does not increase over that transition point the proposed scheme therefore presents a higher-than-average power saving advantage. In the `700HD` scenario the same occurs, but more frequently. This is due to the fact that in this scenario the network nodes have eight times more linecards, so the effect occurs eight times more than in the `150SD` case. A similar effect occurs in the futuristic scenario. But, as the baseline power consumption is much higher than in the first two, the bumps are less pronounced, and are hence imperceptible in the figure.

We now turn to the line in the plot we have not mentioned yet: `3kUHD_ep`. As is well known, current network equipment is not energy proportional [21]. The baseline power (from maintaining the chassis powered on) is very high and is, by a large margin, the main component of router power consumption. In the future it is expected network equipment to increasingly present a more energy proportional profile, so in the plots we also include, for the futuristic scenario `3kUHD`, the situation where all routers are energy-proportional (EP) (`3kUHD_ep`). As can be observed, using the *selective pre-joining* scheme leads to a higher relative gain considering that different starting point in the analysis (i.e., the use of EP routers).

2) Energy consumption model of the hybrid nodes: To quantify the energy savings achieved by introducing optical bypass in an optical IPTV network, in this section we construct a power consumption model of the hybrid node of our solution.

Three factors affect the power consumption of a hybrid node. First, the power consumption of the router. Second, the power consumption of the OXC. Third, the power consumption of the OEO converters (transponders). Note that in this analysis we do not consider the power consumption of other optical equipment that is necessary in an optical network, such as the optical amplifiers, multiplexers and demultiplexers. Previous work [13], [12] has shown that switching equipment and transponders (OEO converters) are the main contributors for power consumption of optical IP networks (responsible for over 97% of total power consumption according to [13]), so we consider switching equipment and OEO converters only. Based on these three variables, we use the following model for the power consumption P of a hybrid node:

$$P = P_R + P_{OEO} + P_{OXC} \tag{4}$$

In Equation 4 P_R is the power consumption of the router, P_{OEO} is the power consumption of the OEO converters (transponders), and P_{OXC} is the power consumption of the optical cross connect. For $P_R + P_{OEO}$ we use the model represented by Equation 1. The power consumption of the OXC is given by Equation 5.

$$P_{OXC} = K_{op} P_{op} \tag{5}$$

In this equation, K_{op} is the number of input/output optical switch ports and P_{op} is the power per input/output switch port. We assume the OXC switching fabric to be based on micro-electro-mechanical systems (MEMS [22]). In a MEMS optical switch, a micro-mirror is used to reflect a light beam. The direction in which the light beam is reflected can be changed by rotating the mirror to different angles, allowing the input light to be connected to any output port. These MEMS have switching times of the order of milliseconds or hundreds of microseconds and for this reason can be used only for slow switching (i.e., circuit switching). For faster switching Semiconductor Optical Amplifiers (SOAs) could be used. But as MEMs consume less power [23], and as the OXC is not to be used for fast switching, MEMS are the option we consider here. We assume 3D-MEMS [24] in particular. The power per input/output switch port of the OXC corresponds to the power consumption for its continuous control, which is equal to 107 mW per input/output port. This value is based on the power consumption of the MEMS controller circuitry of an 80×80 3D-MEMS switch implementation, reported in [24]. We are therefore assuming power consumption is proportional to the number of active input/output ports[6]. The experimental figure and this assumption were considered in previous related work [23], [22] and are also in agreement with studies from other researchers [25], [10].

3) Results: We now analyse how the introduction of optical bypass techniques in the IPTV network translates into energy savings. We consider the same three scenarios as before: `150SD`, `700HD` and `3kUHD`. For the router model we also make the same assumptions. For the first scenario, we assume a router with 4 linecards with 4x1Gbps Ethernet ports each. For the other scenarios we just scale up the model by increasing the number and changing the type of linecards. This implies that

[6] If we assume an on/off behaviour, i.e., a switch consuming its 8.5 W of total power independently of the number of active ports, all results we present change by less than 1%. This stems from the fact that the OXC is the node component with the lowest power consumption by a good margin, in any case.

each wavelength can carry 1Gbps in the first two scenarios, but it scales to 40 Gbps in the third. Note that in this scheme two sets of wavelengths are needed: one for the traffic that optically bypasses the routers, and another for the rest. This is considered in the analysis to calculate the number of active OXC ports. The number of active OEO converters is equal to the number of active ports in the router.

In accordance with the results presented in Figure 5, we assume that 50% of the IPTV traffic optically bypasses the routers, 30% is sent to the router for electronic processing, and 20% of the TV channels are not distributed. Considering this, the power savings for all three scenarios (and the `3kUHD_ep` scenario) are presented in Figure 7. The dashed lines represent the results from using *selective pre-joining* only (the solid lines in Figure 6). The solid lines in the current figure represent the power savings using the optical bypass protocol proposed.

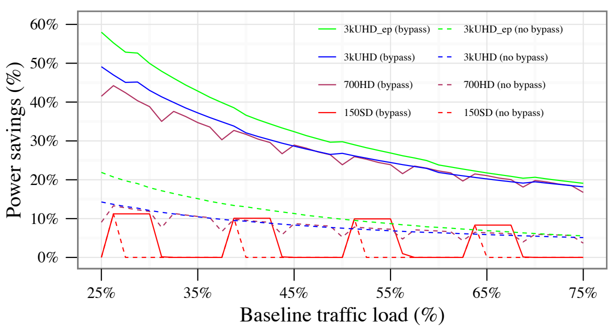

Fig. 7. Power savings achieved by optically bypassing popular TV channels in the network core

When compared with *selective pre-joining* [9], the introduction of optical bypass in the optical IPTV network core increases power savings substantially. At baseline traffic loads of around 30%, the power savings increase from 10% to over 40%, a significant four-fold increase. Considering EP routers, power consumption is halved. We conclude that the use of this technique is very effective in reducing power consumption, including in current IPTV service scenarios (such as `700HD`).

D. Discussion: on the value of electronics

We have just showed how optically switching *popular* IPTV traffic reduces power consumption significantly. How about optically switching *all* IPTV traffic? To answer, we invite the reader to Figure 8. This graph shows the result of optically switching all IPTV traffic in the network core (solid lines), against optically switching only the popular TV channels (dashed lines). As can be observed, by optically switching all IPTV traffic the power savings increase even further. Considering a baseline traffic load of 30%, in most scenarios an additional 15% power saving is achievable by all TV channels bypassing the routers.

So why not moving completely to optics in the future? In a scenario where all IPTV traffic is optically bypassed, to guarantee their availability for IPTV users all TV channels need to be distributed continuously in the network core. This is because OXCs allow slow switching only. The advantage of maintaining the electronic routing option is that, contrary to circuit-switched optical networks, with electronic routing it is

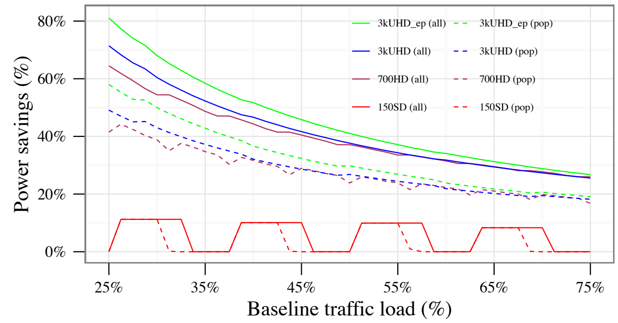

Fig. 8. Power savings achieved by optically bypassing all TV channels in the network core (compared to popular TV channels only)

possible not to distribute all TV channels. This added capability increases bandwidth efficiency. This can be seen in Table III, where we have included the bandwidth savings achieved by not distributing all TV channels. With the increased popularity of narrowcasting services and niche channels, the number of unpopular channels (as per our definition in this paper) may plausibly increase to the several hundreds or thousands in the near future. This trend offers an important argument for the maintenance of electronic routing as an option. A hybrid scheme as the one proposed therefore offers a compromise between *energy* and *resource* efficiency.

VI. Related work

IPTV measurement. With the recent deployment of IPTV networks a number of papers measuring and characterising IPTV traffic has been published [1], [3]. The analysis of real IPTV workloads led to a clearer understanding of how people watch TV and how this impacts the network. The findings from these studies and the analysis of our own dataset offered indications that led to the technique proposed in this paper.

Optics meet electronics. Optical switching techniques such as optical bypass – the technique we propose in this paper for IPTV – have been proposed as an interesting option to reduce the energy footprint of networks [26]. The problem, as explained before, is that due to its coarse granularity, bulk transport in optics can be bandwidth inefficient, especially for bursty traffic. With electronic switching the packets or flows can be processed at a much finer granularity. Smartly combining the strengths of optics and electronics has therefore been considered before [23]. For instance, Huang and Copeland [27] proposed a hybrid routing scheme that can preserve the benefits of optical bypass for large traffic flows and still provide multiplexing gain for small traffic flows. This technique is similar to the one we propose here for IPTV.

Green networking. Since the seminal paper by Gupta and Singh [28] the subject of *green networking* has received considerable attention. Several approaches have been considered to reduce energy consumption in networks, including performing resource consolidation by means of traffic engineering [28] or by putting components to sleep during periods of low traffic activity [29]. The literature in this subject is already substantial, so we refer the interested reader to a more detailed survey on green networking by Bianzino *et al.* [30].

Green IPTV. Recently, caching techniques for reducing energy consumption for time-shifted IPTV systems have been proposed. By considering the particular properties of this type of traffic (which is in several aspects radically different from the live TV broadcasts we consider here), Nencioni *et al.* [7] proposed to cache content on local user storage thereby offloading traffic that would result from subsequent catch-up access. Osman *et al.* [8] also propose a caching strategy to store the most popular programs at nodes closer to the user, considering an IP-over-WDM network. The main differentiating factor of these works against ours is the fact that the IPTV services they consider are VoD or time-shifted broadcasts – not linear TV. Indeed, caching techniques are not suited for this type of service. In our previous work*et al.* [9], we have proposed a scheme that pre-joins only a selection of TV channels, instead of all, to save bandwidth and energy. The protocol we propose in this paper is based on a different technique that goes further in energy efficiency by considering the introduction of optical switching. As a consequence, we significantly improve the energy-efficiency of IPTV networks when compared to [9].

VII. Conclusions

In this paper we considered the introduction of energy-friendly optical technologies to reduce the energy consumption of IPTV distribution networks. We proposed an energy and resource-friendly protocol for the IPTV network core, blending electronic routing with all-optical switching. The main idea is to optically switch popular TV channels. This IPTV traffic bypasses the routers and therefore does not require any electronic processing (it is switched at the optical layer). The rest of the channels are sent to the routers for electronic processing (to be switched at the IP layer). By analysing a large dataset from an IPTV operator, we observed that with the proposed protocol it is possible to switch 50% of the IPTV traffic all-optically. The energy savings obtained from optically bypassing this traffic are substantial, reaching power savings of over 40% under normal load conditions. The scheme is also bandwidth efficient as channels without viewers are not distributed.

Acknowledgment

We would like to thank the anonymous reviewers for their feedback. This project has received funding from the European Union's FP7 research and Innovation programme under grant agreement No FP7-607109 (SEGRID) and from the Portuguese State through funds assigned by Fundação para a Ciência e a Tecnologia (FCT) to LaSIGE Research Unit, ref. UID/CEC/00408/2013.

References

[1] M. Cha et al. Watching television over an IP network. In *IMC*, Vouliagmeni, Greece, October 2008.

[2] Digital TV Research. Global IPTV Forecasts report. Technical report, Digital TV Research, 2014.

[3] T. Qiu, Z. Ge, S. Lee, J. Wang, J. Xu, and Q. Zhao. Modeling user activities in a large IPTV system. In *IMC*, Chicago, IL, November 2009.

[4] F. Legendre et al. Narrowcasting: an empirical performance evaluation study. In *CHANTS*, San Francisco, CA, September 2008.

[5] IPCC. Climate Change 2014: Synthesis Report. Technical report, The Climate Group, 2014.

[6] M. Webb. SMART 2020: Enabling the low carbon economy in the information age. Technical report, M. Webb, 2008.

[7] Gianfranco Nencioni et al. Understanding and decreasing the network footprint of catch-up tv. In *Proceedings of the 22Nd International Conference on World Wide Web*, WWW '13, 2013.

[8] N.I. Osman et al. Energy-efficient future high-definition tv. *Lightwave Technology, Journal of*, 32(13), July 2014.

[9] F. M. V. Ramos et al. Reducing energy consumption in IPTV networks by selective pre-joining of channels. In *SIGCOMM workshop on green networking*, New Delhi, India, August 2010.

[10] J. Baliga et al. Energy consumption in optical IP networks. *Journal of Lightwave Technology*, 27(13):2391–2403, 2009.

[11] C. Fraleigh et al. Packet-level traffic measurements from the Sprint IP backbone. *IEEE Network*, 17(6):6–16, 2003.

[12] Y. Zhang et al. Energy efficiency in telecom optical networks. *IEEE Communications Surveys and Tutorials*, 12(4):441–458, 2010.

[13] G. Shen and R. S. Tucker. Energy-minimized design for IP over WDM networks. *IEEE/OSA Journal of Optical Communications and Networking*, 1(1):176–186, 2009.

[14] C.-H. Gan, P. Lin, and C.-M. Chen. A novel prebuffering scheme for IPTV service. *Computer Networks*, 53(11):1956–1966, 2009.

[15] C. Y. Lee, C. K. Hong, and K. Y. Lee. Reducing channel zapping time in IPTV based on user's channel selection behaviors. *IEEE Transactions on Broadcasting*, 56(3):321–330, 2010.

[16] T. Qiu et al. Modeling channel popularity dynamics in a large IPTV system. In *SIGMETRICS*, Seattle, WA, June 2009.

[17] P. Mahadevan et al. A power benchmarking framework for network devices. In *IFIP*, Buenos Aires, Argentina, June 2009.

[18] V. Sivaraman et al. Profiling per-packet and per-byte energy consumption in the NetFPGA Gigabit router. In *INFOCOM Workshop on Green Communications and Networking*, Shanghai, China, April 2011.

[19] Alcatel. Alcatel-Lucent WaveStar OLS 1.6T product specification. http://tinyurl.com/AlcatelWavestar. [Online; accessed 01-12-2015].

[20] W. Hou et al. Green multicast grooming based on optical bypass technology. *Optical Fiber Technology*, 17(2):111–119, 2011.

[21] L. A. Barroso and U. Holzle. The case for energy-proportional computing. *IEEE Computer*, 40(12):33–37, 2007.

[22] M. Fiorani, M. Casoni, and S. Aleksic. Performance and power consumption analysis of a hybrid optical core node. *IEEE/OSA Journal of Optical Communications and Networking*, 3(6):502–513, 2011.

[23] S. Aleksic. Analysis of power consumption in future high-capacity network nodes. *IEEE/OSA Journal of Optical Communications and Networking*, 1(3):245–258, 2009.

[24] M. Yano, F. Yamagishi, and T. Tsuda. Optical MEMS for photonic switching-compact and stable optical crossconnect switches for simple, fast, and flexible wavelength applications in recent photonic networks. *IEEE Journal of Selected Topics in Quantum Electronics*, 11(2):383–394, 2005.

[25] R. S. Tucker. The role of optics and electronics in high-capacity routers. *Journal of Lightwave Technology*, 24(12):4655–4673, 2006.

[26] K. Hinton et al. The future Internet – an energy consumption perspective. In *OECC*, Hong Kong, China, July 2009.

[27] H. Huang and J. A. Copeland. Optical networks with hybrid routing. *IEEE Journal on Selected Areas in Communications*, 21(7):1063–1070, 2003.

[28] M. Gupta and S. Singh. Greening of the Internet. In *SIGCOMM*, Karlsruhe, Germany, August 2003.

[29] S. Nedevschi et al. Reducing network energy consumption via sleeping and rate-adaptation. In *NSDI*, San Francisco, CA, April 2008.

[30] A. P. Bianzino et al. A survey of green networking research. *IEEE Communications Surveys and Tutorials*, 14(1), 2012.