

Energy-Efficient VoD Content Delivery and Replication in Integrated Metro/Access Networks

Marco Savi, Giacomo Verticale, Massimo Tornatore, Achille Pattavina

Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Piazza Leonardo da Vinci 32, Milano, Italy
firstname.lastname@polimi.it

Abstract—Today’s growth in the demand for access bandwidth is driven by the success of the Video-on-Demand (VoD) bandwidth-consuming service. At the current pace at which network operators increase the end users’ access bandwidth, and with the current network infrastructure, a large amount of video traffic is expected to flood the core/metro segments of the network in the near future, with the consequent risk of congestion and network disruption. There is a growing body of research studying the migration of content towards the users. Further, the current trend towards the integration of metro and access segments of the network makes it possible to deploy Metro Servers (MSes) that may serve video content directly from the novel integrated metro/access segment to keep the VoD traffic as local as possible. This paper investigates a potential risk of this solution, which is the increase in the overall network energy consumption. First, we identify a detailed power model for network equipment and MSes, accounting for fixed and load-proportional contributions. Then, we define a novel strategy for controlling whether to switch MSes and network interfaces on and off so as to strike a balance between the energy consumption for content transport through the network and the energy consumption for processing and storage in the MSes. By means of simulations and taking into account real values for the equipment power consumption, we show that our strategy is effective in providing the least energy consumption for any given traffic load.

I. INTRODUCTION

Internet fixed and mobile traffic is constantly increasing [1]. There are various causes behind such sustained traffic growth. Among them, is the adoption of broadband video streaming services, such as Video-on-Demand (VoD), that is increasing year by year due to the rising user engagement with innovative multimedia services. In few years media streaming and especially video traffic will be dominant in the network infrastructure. Telecom operators are making huge efforts to provide the users with higher access rates by deploying innovative fixed and mobile network access technologies, such as fiber-to-the-home (FTTH), fiber-to-the-building (FTTB) or Long Term Evolution (LTE).

Telecommunication networks can be roughly divided in three different and interconnected segments: the access, the metro/aggregation and the core segment. In the current scenario, video contents are usually stored in centralized servers and/or data centers placed in the core network. The expected video-traffic explosion will significantly contribute to flood the core network, that might become the new bottleneck of the overall infrastructure, severely compromising the quality of all Internet services.

A possible solution to avoid such scenario is to migrate video content servers closer to the users, possibly in the metro and/or access segments. With the deployment of a system of distributed video servers (that we call Metro Servers (MSes)) in these segments of the network, a high amount of VoD traffic can be kept locally. The core network will benefit from this offloading and the users will experience a better quality of experience (QoE), since the content delivery latency will be consistently decreased.

Even though the advantages in terms of offloading and latency are evident, the use of video servers in the metro/access area must be carefully studied and monitored to avoid an indiscriminate increase of the energy consumption of this network segments. In some existing studies [2][3][4][5], it has already been shown that the deployment of distributed VoD servers with respect to a centralized solution leads to two contrasting effects. On one side, the overall *transport energy consumption*, i.e., the energy needed to deliver the video streaming through the network, is reduced, since some of the video contents are now delivered to the users from a closer location. On the other side, the overall energy needed for storage and data processing (i.e., the *caching energy consumption*) of the video content is increased due to the high amount of video servers which must be placed and powered on in the metro/access networks. Therefore, according to traffic conditions, it might be more energy efficient to switch MSes on or off, and, to the best of our knowledge, no study has been performed to investigate how to manage the energy state of MSes in presence of traffic requests dynamically changing.

In this work, starting from an investigation about the energy consumption of VoD delivery in presence of MSes, we will propose and evaluate novel strategies for coordinated on/off switching of the network interfaces and MSes depending on traffic load conditions during the day. These strategies will allow a VoD operator to save a high amount of energy. Specifically, we will consider an integrated metro/access network architecture, where the metro and access segments are consolidated in a single domain and where fixed and mobile access technologies share the same infrastructure [6].

The reminder of the paper is organized as follows: in Section II we provide an overview of related works, while in Section III we describe our reference integrated network architecture. In Section IV we describe the power models adopted for the network equipment, and in Section V we define our strategies

for coordinated switching on and off the MSes to save energy. Finally, in Section VI we show illustrative numerical examples based on simulations to quantify the gain of our approach. Section VII concludes the paper.

II. RELATED WORKS

There are some existing studies dealing with energy consumption and energy efficiency evaluation related to the deployment of a system of distributed VoD servers closer to the users. Ref. [5] was the first work to investigate the existence of a tradeoff between the energy needed for transport and the energy needed for storage of video contents. Ref. [7] defines an Integer Linear Program (ILP) for caches placement closer to the user while minimizing the overall energy consumption. Both these works assume that contents are replicated closer to the users, but servers are still placed only in the core network, while in our work we consider placing content even closer to the users (i.e., in the metro/access network) to traffic offload the core network. Ref. [8] proposes an ILP for energy-efficient content distribution: in this case, all the backbone nodes feature some caching capabilities and the model identifies the most energy-efficient solution, which allows to save energy switching off some links and some caches. Even though this work (as Ref. [5] and Ref. [7]) refers to content distribution in the core network, it is very relevant to our work as it deals with the strategy of switching off some network equipment.

Ref. [3] defines five strategic locations for VoD content placement spanning from the access to the core network and comparatively analyses transport and storage energy consumption variation while storing contents at these locations. Authors also define some thresholds for the VoD request arrival rate which are useful to understand when it is energy efficient to switch on/off content servers for the different locations, considering an energy proportional model [9] for all the network equipment. In our approach for the first time we analyse content distribution in a specific integrated metro/access network architecture providing a coordinated on/off strategy, that can be applied in a dynamic network scenario. We also overcome the assumption of a wholly proportional energy consumption model, which we consider to be inaccurate in most of today's network equipment, where energy consumption is only slightly dependent by the carried load [10], as we will point out in Section IV.

III. THE NETWORK ARCHITECTURE

Figure 1 shows the integrated metro/access architecture considered in this paper, which consists in a tree-topology network infrastructure consolidating the metro and access segments of the network. All the nodes are active, i.e., switches, as opposed to the passive devices such as splitters used in Passive Optical Networks (PONs). For a comparison between active and passive architectures, the reader is referred to [11].

We suppose that a Content Delivery Server (CDS) storing all the contents is placed in the core network, then we add the possibility of placing Metro Servers (MSes) at different

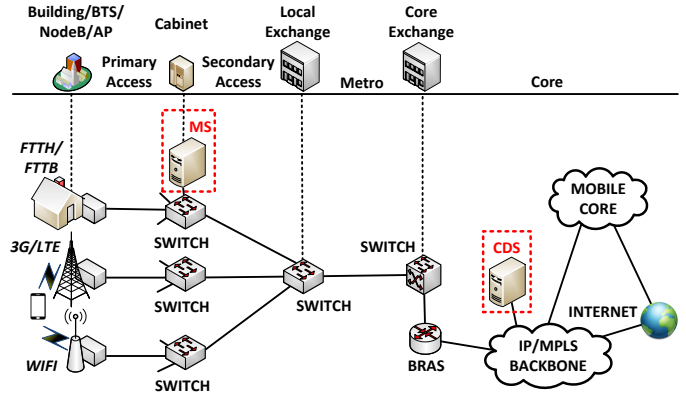


Fig. 1. The integrated metro/access network architecture with fixed and mobile infrastructure sharing

positions in the metro/access network. The network architecture allows resource sharing between fixed access technologies (e.g., FTTH, FTTB, ADSL or WiFi) and mobile access technologies (e.g., 3G or LTE). Fixed and mobile traffic is delivered from/to the core network through the switches placed in the metro/access tree network topology. In this architecture the MSes allow heterogeneous users to retrieve VoD content from a closer locations, enabling traffic offloading of the core network and reducing latency for content retrieval. Note that metro and access consolidation as well as an infrastructure sharing between fixed and mobile access technologies are two of the most promising ideas for the deployment of cost- and energy-efficient next generation networks [6].

It is important to stress that Fig. 1 is just an example of the reference architecture, in which the metro/access is a two stage tree and the MSes are placed in the street cabinets. Depending on the geographical and hierarchical structure of the actual network, the tree topology can vary the number of the stages and/or the position of the MSes, or evolve towards more complex meshed or ring topologies. The main conceptual contribution of this paper could be easily generalized also to those cases. In the reminder of the paper we will always refer to this tree-based network architecture for the energy consumption evaluation of the VoD service.

IV. SWITCH AND METRO SERVER POWER MODELS

In this Section we describe the power consumption models for the network equipment used in the VoD service deployment. The two main elements, according to the network architecture depicted in Section III, are the switches, which contribute to the *transport power* consumption, and the Metro Servers, which contribute to the *caching power* consumption. Especially, we investigate the possibility of switching on only the minimum number of network interfaces to guarantee the VoD service provisioning and to save part of the transport power consumption.

A. Switch power consumption model

In our model, we split the power consumption of the switch in two different contributions, similarly to [12]:




Fig. 2. The switching and the single link transmission power consumption curves competing in the switch power model

- 1) The switching power consumption P_{sw}
- 2) The transmission power consumption P_{tr}

The *switching power* consumption P_{sw} is related to all the operations needed to forward traffic from an incoming interface to an outgoing interface. For example, the power consumption of the backplane/switching fabrics of the switch is considered in this power consumption contribution. This contribution is highly traffic dependent (Fig. 2a).

The *transmission power* consumption P_{tr} is related to the transmission devices (i.e., network interfaces) which are active and transmit data. In this case, once a network interface is active, it consumes a fixed amount of power, that we call P_{int} , independently from the traffic load transmitted by that interface. Note that when traffic grows it can be necessary to power on more than one network interface (each one of capacity C_{int}) over a link to serve all the VoD requests (in this way we can increase the bandwidth C_{link} of a link and avoid blocking of video requests). On the contrary, when traffic decreases, one or more interfaces can be switched off. So, for each link, the power consumption P_{link} can be represented as a step function of the traffic load to be transmitted (Fig. 2b).

Note that we suppose that switches are always powered on as they have to support also other Internet services and not only VoD delivery. Therefore, one network interface for each link will be always active to provide connectivity to all the nodes in the integrated metro/access network and its power consumption is not considered in the overall VoD transmission power consumption. Based on these remarks, we can write the formula for P_{switch} , i.e., the overall power consumption of the switch to support the VoD service:

$$\begin{aligned}
 P_{switch}(\vartheta) &= P_{sw}(\vartheta) + P_{tr}(\vartheta) \\
 &= P_{sw}(\vartheta) + \sum_{i \in V^+} P_{link}^i(\vartheta_i) \\
 &= P_{sw}(\vartheta) + \sum_{i \in V^+} (n_i(\vartheta_i) - 1) \cdot P_{int}
 \end{aligned}$$




Fig. 3. The MS power model

where ϑ is the load processed by the switch expressed in bit/s, V^+ is the set of the outgoing links, ϑ_i is the load transmitted over the link $i \in V^+$ and $n_i(\vartheta_i)$ is the number of active ports for the transmission over the link i , which depends on the traffic load to be transmitted over that link. We call P_{link}^i the overall power consumed by the switch for the transmission over the link i . Note also that $\vartheta = \sum_{i \in V^+} \vartheta_i$, because we suppose that all the traffic is forwarded and transmitted by the switch, and that the capacity C_{link}^i of each link i can be expressed as $C_{link}^i(\vartheta_i) = n_i(\vartheta_i) \cdot C_{int}$.

This model assumes that the switch is able to efficiently aggregate traffic to minimize the number of used network interfaces.

B. Metro Server power consumption model

As for the switch power model, we assume that the power consumption of an MS can be modelled by two different contribution:

- 1) The idle power consumption P_{idle}
- 2) The processing power consumption P_{proc}

The *idle power* consumption P_{idle} is the load-independent power consumption of the MS. As reported in [10], the load-independent amount of energy consumed by a server is dominant, and it is mainly due to the idle power consumed by the CPUs and to the storage disks.

On the contrary, the *processing power* consumption P_{proc} is the load-dependent part of the MS power consumption, and it is proportional to the processed load. This contribution is mainly due to CPUs load-dependent power consumption. Note that today's MS energy consumption is far from being proportional, even though many proposals in this direction have been investigated [9]. Hence, we can write the formula for P_{ms} , i.e., the power consumption of the MS (Fig. 3):

$$P_{ms}(\vartheta) = P_{idle} + P_{proc}(\vartheta)$$

Concerning the CDS, which is supposed to be always powered on and to store all the content as remarked in Section III, we consider an energy proportional behavior. This assumption holds as in the centralized data centers economies of scale can be easily exploited to minimize the P_{idle} of the equipment. So, we can model P_{cds} , i.e., the power consumption of the CDS, in the following way:

$$P_{cds}(\vartheta) = P_{proc}(\vartheta)$$

V. AN ENERGY-EFFICIENT COORDINATED STRATEGY TO CONTROL THE POWER STATUS OF DISTRIBUTED MSES

In Section IV we described our models for the energy consumption of switches and MSes, that yield two main contribution for the overall VoD service energy consumption: the *transport energy* and the *caching energy*. In this Section, we propose to address the tradeoff between these two energy contributions when operating VoD services over the network architecture described in Section III. Specifically, our aim is to minimize energy consumption by properly switching on and off the MSes during the day according to the change of traffic conditions. Considering our power models for switches and MSes/CDS, the best strategy to power on or off the MSes depends 1) on the normalized *diurnal VoD request rate distribution*, that we define as $\lambda_{norm}(t)$ (i.e., when during the day there are traffic peaks), and 2) on the *aggregated VoD traffic*, ρ , which represents the total throughput of the VoD traffic served to the end users averaged over the day, and grows as the number of the daily VoD requests grows. In other words, when the network is expected to be flooded by a high number of VoD requests, it will be more convenient to serve them at the MSes in order to save a high amount of transport energy. In fact, in this case the switches at higher levels of the metro/access network and of the core network will have to process a lower amount of VoD traffic. Thus some network interfaces, as described in Section IV, could be switched off, leading to relevant savings in P_{switch} , as the traffic is kept local. On the contrary, when the request rate is low, serving VoD content directly at the CDS can lead to energy savings as the MSes can be switched off to save their P_{idle} power consumption.

In the reminder of this Section, we first define a strategy to plan different on/off periods (*on/off patterns*) for the MSes according to the diurnal VoD request rate distribution, trying to capture different normalized traffic peaks during the day.

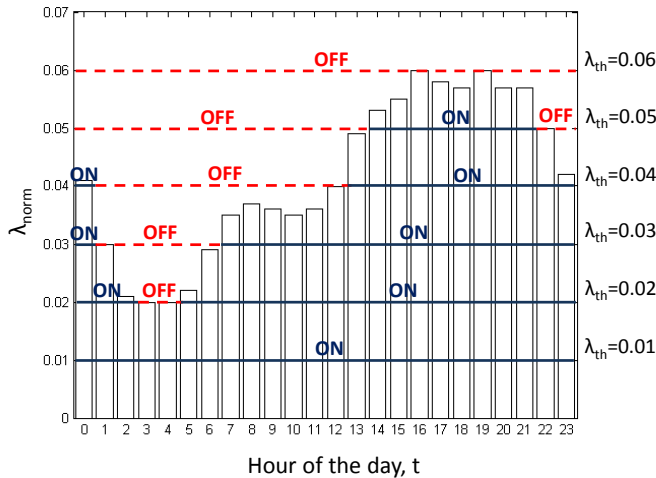


Fig. 4. An example of diurnal VoD request rate distribution together with a graphical illustration of different on/off patterns, obtained by the application of the on/off strategy for different values of λ_{th}

TABLE I
SUMMARY OF THE ON/OFF PATTERNS OBTAINED BY DIFFERENT VALUES FOR λ_{th} , AS DEPICTED IN FIG. 4

λ_{th}	On/off pattern	
	On hours	Off hours
0.01	All the day	—
0.02	05:00 am - 02:59 am	03:00 am - 04:59 am
0.03	07:00 am - 12:59 am	01:00 am - 06:59 am
0.04	01:00 pm - 12:59 am	01:00 am - 12:59 pm
0.05	02:00 pm - 09:59 pm	10:00 pm - 01:59 pm
0.06	—	All the day

Then, we define which is the most energy-efficient on/off pattern when a specific value for the aggregated VoD traffic ρ is considered.

A. MSes on/off patterns planning

A VoD operator can plan different daily on/off patterns for the MSes by analyzing the diurnal VoD request rate distribution, that we suppose can be estimated in an accurate way by the VoD operator itself. By defining the parameter λ_{th} as the *VoD request rate threshold*, we can set the following *on/off strategy*.

Definition 1 (On/off strategy). *Considering the diurnal VoD request rate distribution over time $\lambda_{norm}(t)$ and a VoD request rate threshold λ_{th} , the MSes are switched on or switched off depending on λ_{th} according to the following rule:*

- If $\lambda_{norm}(t) \leq \lambda_{th}$, then at time t the MSes are off
- If $\lambda_{norm}(t) > \lambda_{th}$, then at time t the MSes are on

The output of the on/off strategy is a daily *on/off pattern*. If we consider different values for λ_{th} , we can obtain different on/off patterns, each one respective to a specific value for λ_{th} . For illustrative purposes, let us consider the diurnal VoD request rate distribution estimated in [13] and reported in Fig. 4, where the request rate varies hourly. In this specific scenario, we apply the on/off strategy considering some values for λ_{th} , obtaining the respective on/off patterns, as depicted in Fig. 4 and further reported in Table I. We choose these values because they allow to capture the main trends of the diurnal VoD request rate distribution. Note that, in general, the more values of λ_{th} are considered, the more different on/off patterns can be explored.

B. Definition of the best on/off pattern

Once different on/off patterns have been defined, the next step for a VoD operator consists in exploring and estimating which is the *best on/off pattern*, i.e., which on/off pattern




Fig. 5. The Best Energy Pattern Algorithm (BEPA)




Fig. 6. A 64 NT metro/access network with three splitting stages and a MS placed in the second splitting stage

leads to the minimum daily energy consumption, when a given aggregated VoD traffic ρ is considered. Note that as ρ changes, e.g. considering different days of the week and/or different periods of the year, the best on/off pattern changes accordingly. Our proposed approach allows a VoD operator to always select the best on/off pattern for any value of ρ , by evaluating an *energy consumption estimation* using the models shown in Section IV. Figure 5 shows our overall strategy, that we call *Best Energy Pattern Algorithm* (BEPA), for the best on/off pattern definition. In Section VI we will apply our Best Energy Pattern Algorithm to simulated aggregated VoD traffic considering today's network equipment power consumption.

VI. NUMERICAL RESULTS

A. Overview of the simulator

We developed a discrete-event simulator using Matlab. A detailed description of our adopted VoD content and traffic modelling can be found in [11]. In summary, the video requests probability distribution can be modelled by a Poisson distribution, as in [3], where for every hour of the day the request rate is obtained by the diurnal VoD request rate distribution as shown in Fig. 4. Our simulator reproduces, for each video request, the requested bandwidth, its duration and the assignment to a MS or to the CDS. We consider the network scenario in Fig. 6: we simulate an integrated metro/access active tree topology, as described in Section III, with a per-stage splitting ratio of 1:4 and three splitting stages, leading to an overall splitting ratio of 1:64 (64 network terminations, 64 NT). As depicted, we consider the placement of a single MS in the second splitting stage of the network. We assume that all the video contents are replicated in the MS and are encoded considering four different streaming bitrates. The parameters we used in our simulation are defined in [3][10] and are reported in Table II.

TABLE II
SIMULATION PARAMETERS

Parameter	Value [3][10]
P_{sw}	$3.31 \cdot 10^{-9} \frac{W}{bit/s}$
P_{port}	15W
P_{proc}	$2.16 \cdot 10^{-8} \frac{W}{bit/s}$
P_{idle}	156W
C_{port}	1 Gbit/s
# Average core network hops	10 hops




Fig. 7. Average daily energy consumption for different values of ρ when the MS is always on, always off and applying the Best Energy Pattern Algorithm

B. Gain of the Best Energy Pattern Algorithm

Figure 7 shows the daily energy consumption of the VoD service as a function of the aggregated VoD traffic ρ , considering 1) the MS always turned on (MS ON curve), 2) the MS always turned off (MS OFF curve) and 3) the best on/off patterns chosen for each value of ρ applying the Best Energy Pattern Algorithm (BEPA curve). Figure 7 shows that when $\rho < 1$ Gbit/s, the BEPA and the MS ON curves overlap, i.e., for low traffic conditions it is always more energy efficient to serve VoD contents from the CDS and turning off the MS, since the overall caching energy consumption while keeping the MS always on is much higher than the overall transport energy consumption while keeping the MS always off. Conversely, when $\rho > 1.6$ Gbit/s, the BEPA and the MS OFF curves overlap, as it is always more energy efficient to keep on the MS for the whole day and to serve the VoD requests closer to the users. In this case, the overall transport energy consumption while keeping the MS always off is much higher than the overall caching energy consumption while keeping the MS always turned on.

When we consider ρ values that lead to a balance between the caching energy consumption and the transport energy consumption, the Best Energy Pattern Algorithm allows to save energy by properly switching on/off the MSes during the day. This happens around the intersection point between the MS ON curve and the MS OFF curve. In this configuration, it happens when $\rho \simeq 1.3$ Gbit/s. Note that the intersection between MS ON and MS OFF moves towards higher values of ρ when the overall transport energy consumption becomes less relevant than the overall caching energy consumption




Fig. 8. Average daily energy consumption for different values of λ_{th} considering different values of ρ . The matching between λ_{th} and the respective on/off pattern is reported in Table I and Fig. 4

(for example, if P_{sw} and/or P_{port} decrease). Vice versa, the intersection would happen for lower values of ρ .

Figure 8 shows the daily energy consumption for different values of λ_{th} , while keeping the aggregate VoD traffic ρ constant, around the value of $\rho \simeq 1.3$ Gbit/s. When $\rho = 1.3$ Gbit/s, we can see that $\lambda_{th} = 0.03$ leads to the minimum energy consumption. According to the matching between λ_{th} and the on/off pattern reported in Table I, this means that turning the MS off from 01:00 am to 06:59 am, under our assumptions, is the best solution, which leads to about 15% of energy savings with respect to always powering the MS on ($\lambda_{th} = 0.01$) and about 25% with respect to always keeping it off ($\lambda_{th} = 0.06$). Considering $\rho = 1$ Gbit/s, the maximum energy savings are obtained for $\lambda_{th} = 0.05$, i.e., when the MS is turned on from 02:00 pm to 09:59 pm. In this case, this on/off pattern allows to save about 15% of energy in comparison with always keeping the MS off and about 35% in comparison with always powering it on. Then, in case of $\rho = 1.6$ Gbit/s, we can see that the best solution is again turning the MS off from 01:00 am to 06:59 am ($\lambda_{th} = 0.03$). Considering this traffic intensity, energy savings of about 10% can be achieved with respect to always powering the MS on and of about 25% with respect to always powering it off.

VII. CONCLUSION

In this work we explored the energy consumption for VoD service delivery when distributed video servers (MSes) are deployed in an integrated metro/access network supporting both fixed and mobile services. We defined a power consumption model for the equipment involved in the video content delivery, i.e. the switches, the MSes and the centralized Content Delivery Server. Our models capture the load-dependent and the load-independent behaviour of the energy consumption of the network equipment and take into account the possibility of switching on/off some network interfaces according to VoD traffic intensity carried over the links. We then investigated an effective strategy to define daily on/off patterns for the MSes depending on the VoD request rate variation during the day

and we provided a tool to always choose the best on/off pattern depending on traffic requirements. We investigated by means of simulations the effectiveness of our approach considering real equipment power consumption.

Several issues still remain open for future research. For example, more complex on/off strategies for the MSes can be explored as well as more refined power models for the equipment. Moreover, real request data could be used to improve simulation results. Currently, we are also investigating the advantages that the Software Defined Networking (SDN) paradigm could stimulate from an energy efficiency perspective, by including in the metro/access architecture a centralized controller and SDN-enabled switches.

ACKNOWLEDGMENT

The research leading to these results has received funding from the European Community Seventh Framework Programme FP7/2013-2015 under grant agreement no. 317762 COMBO project.

REFERENCES

- [1] "Cisco visual networking index: Global mobile data traffic forecast update, 2012-2017," February 2013.
- [2] K. Hinton, J. Baliga, M. Feng, R. Ayre, and R. Tucker, "Power consumption and energy efficiency in the internet," *Network, IEEE*, vol. 25, no. 2, pp. 6–12, March 2011.
- [3] C. Jayasundara, A. Nirmalathas, E. Wong, and C. A. Chan, "Improving energy efficiency of video on demand services," *Optical Communications and Networking, IEEE/OSA Journal of*, vol. 3, no. 11, pp. 870–880, November 2011.
- [4] E. Di Pascale, D. Payne, and M. Ruffini, "Bandwidth and energy savings of locality-aware P2P content distribution in next-generation PONs," in *Optical Network Design and Modeling (ONDM), 2012 16th International Conference on*, 2012, pp. 1–6.
- [5] J. Baliga, R. Ayre, K. Hinton, and R. Tucker, "Architectures for energy-efficient IPTV networks," in *Optical Fiber Communication - includes post deadline papers, 2009. OFC 2009. Conference on*, March 2009, pp. 1–3.
- [6] "European FP7 project COMBO," <http://www.ict-combo.eu>, 2013-2015.
- [7] U. Mandal, C. Lange, A. Gladisch, P. Chowdhury, and B. Mukherjee, "Energy-efficient content distribution over telecom network infrastructure," in *Transparent Optical Networks (ICTON), 2011 13th International Conference on*, June 2011, pp. 1–4.
- [8] J. Araujo, F. Giroire, Y. Liu, R. Modrzejewski, and J. Moulrierac, "Energy efficient content distribution," in *Communications (ICC), 2013 IEEE International Conference on*, June 2013, pp. 4233–4238.
- [9] L. Barroso and U. Holzle, "The case for energy-proportional computing," *Computer*, vol. 40, no. 12, pp. 33–37, Dec 2007.
- [10] D. Tsirogiannis, S. Harizopoulos, and M. A. Shah, "Analyzing the energy efficiency of a database server," in *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data*, ser. SIGMOD '10. New York, NY, USA: ACM, 2010, pp. 231–242. [Online]. Available: <http://doi.acm.org/10.1145/1807167.1807194>
- [11] R. Fratini, M. Savi, G. Verticale, and M. Tornatore, "Using replicated video servers for VoD traffic offloading in integrated Metro/Access networks," in *IEEE ICC 2014 - Optical Networks and Systems (ICC'14 ONS)*, Sydney, Australia, Jun. 2014.
- [12] P. Mahadevan, P. Sharma, S. Banerjee, and P. Ranganathan, "A power benchmarking framework for network devices," in *NETWORKING 2009*. Springer, 2009, pp. 795–808.
- [13] W. Tang, Y. Fu, L. Cherkasova, and A. Vahdat, "Medisyn: A synthetic streaming media service workload generator," in *Proceedings of the 13th International Workshop on Network and Operating Systems Support for Digital Audio and Video*, ser. NOSSDAV '03. New York, NY, USA: ACM, 2003, pp. 12–21. [Online]. Available: <http://doi.acm.org/10.1145/776322.776327>