

Energy Efficient Content Distribution in an ISP Network

Remigiusz Modrzejewski, Luca Chiaraviglio, Issam Tahiri, Frederic Giroire, Esther Le Rouzic, Edoardo Bonetto,
Francesco Musumeci, Roberto Gonzalez, Carmen Guerrero

Abstract—We study the problem of reducing power consumption in an Internet Service Provider (ISP) network by designing the content distribution infrastructure managed by the operator. We propose an algorithm to optimally decide where to cache the content inside the ISP network. We evaluate our solution over two case studies driven by operators feedback. Results show that the energy-efficient design of the content infrastructure brings substantial savings, both in terms of energy and in terms of bandwidth required at the peering point of the operator. Moreover, we study the impact of the content characteristics and the power consumption models. Finally, we derive some insights for the design of future energy-aware networks.

I. INTRODUCTION

The electricity consumption of the Information and Communication Technology (ICT) sector represents today almost 2% of the world electricity [1], having observed an annual increase of 10% from 2007 to 2012. In this context, data centers and backbone networks will experience the highest energy consumption growth rates in the forthcoming years [2], due to the increase of traffic, especially for multimedia content. As an example, Cisco estimates that the sum of all forms of videos will represent 86% of the global consumer traffic by 2016 [3]. In order to mitigate this trend, different solutions have been proposed in the literature for the design and the management of energy-efficient backbone networks (see [4] for an overview).

Recently, the problem of reducing power consumption in a backbone network by moving the contents accessed by users has attracted attention of the research community. In particular, in [5] we have studied the problem of reducing power consumption of an Internet Service Provider (ISP) and a Content Provider (CP) jointly, showing that considerable

The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n. 257740 (Network of Excellence "TREND"). Special thanks to Ian Pye of CloudFlare Inc. for providing information on power consumption of in-network caches.

R. Modrzejewski, I. Tahiri and F. Giroire are with COATI, INRIA, I3S, CNRS, Université de Nice Sophia, Sophia Antipolis, France, (e-mail: {firstname.lastname}@inria.fr).

L. Chiaraviglio is with CNIT and University of Rome "Sapienza", Rome, Italy (e-mail: luca.chiaraviglio@diet.uniroma1.it).

E. Le Rouzic is with Orange Labs, Networks and Carriers, Lannion, France (e-mail: esther.lerouzic@orange.com).

E. Bonetto is with the Department of Electronics and Telecommunications, Politecnico di Torino, Torino, Italy (e-mail: edoardo.bonetto@polito.it).

F. Musumeci is with CNIT and Politecnico di Milano, Department of Electronics and Information, Milano, Italy (e-mail: fmusumeci@elet.polimi.it).

R. Gonzalez and C. Guerrero are with Universidad Carlos III de Madrid, Madrid, Spain (e-mail: {rgonza1@it.uc3m.es, carmen.guerrero@uc3m.es}).

energy savings can be obtained when the CP and the ISP cooperate to minimize the total power consumption. In [6] authors propose an architecture based on Content Centric Networking (CCN) to reduce the power consumption. Additionally, in [7] an architecture based on home gateways forming a distributed data center infrastructure managed by the ISP is proposed and evaluated. Finally, the energy trade-offs of an architecture based on immersive video centric services are evaluated in [8]. All these works prove that a huge amount of power is saved when the ISP takes control of the content and caches it considering the energy consumed to move the information across the network.

In this work we study the problem of reducing power consumption in an ISP network by considering the design of a content distribution infrastructure managed by the ISP. Our aim is to study where to cache content inside the network in order to reduce the overall power consumption of the system composed by the network elements and the installed storage. In current ISP networks a huge amount of traffic is exchanged between the users and the data centers owned by large CPs, such as Google, Yahoo, Amazon, and Limelight. Normally, the data centers of large CPs are located close to the peering points of the ISP [9]. Therefore, the traffic originated from the data centers has to traverse a number of hops in the ISP network before reaching the users. We therefore study the optimal content caching inside the ISP, rather than sending the content from the data centers to the users. In our scenario, we consider a hierarchical logical topology composed of different levels (e.g. core, metro, and access), and we optimize the energy consumption by choosing the best level where to put each content.

The benefits of the energy-efficient design of content distribution architectures inside the ISP are multiple. First of all, it is possible to reduce jointly the electricity costs of the storage and of the network, as their energy is explicitly taken into account during the design phase. Secondly, the ISP reduces the amount of traffic that is exchanged across the network. This in turn may decrease the maintenance costs incurred by the ISP, since network elements are upgraded less frequently and less new switching devices need to be installed. Third, the monetary costs for sending/receiving information from outside the network are also reduced, since less bandwidth is required.

The closest papers to our work are [10]–[12]. In [10] authors detail an analytical model for caching considering the cost for transporting information and the cost for storing the content. However, the model is derived for a simple scenario (a metro network), with at most three levels in the topology as possible

locations to cache content. Moreover, the evaluation of savings in terms of energy is not performed. In [11] authors propose a model for caching that integrates energy costs. The evaluation is performed considering five possible levels for caching. Finally, in [12] an ILP model and two simple heuristics for the energy-efficient content distribution are detailed. However, authors do not consider the energy consumed for sending the content to the possible locations inside the network and a limited number of levels is also assumed.

In contrast to previous work, in this paper we go one step further by: a) defining a model with a generic number of levels and not only restricted to specific values or specific segments of the network, b) proposing an optimal algorithm to decide where to cache the content and compute the total energy consumption, c) evaluating the results over two case studies. Moreover, we consider the impact of the topology properties on the content caching, and we derive some insights for the design of future energy-aware networks.

The rest of the paper is organized as follows. In Sec. II we describe the problem. Sec. III details the algorithm we propose to solve the problem. Results are presented in Sec. IV. Finally, Sec. V concludes our work.

II. PROBLEM DESCRIPTION

We assume that the network is organized in a hierarchical structure composed of different levels. In particular, we assume a tree-like network to represent the collection of paths between each user and the Internet peering node. Nodes are grouped according to a hierarchy, and each level of the tree corresponds to a different level of the hierarchy. The content data is delivered towards the clients following a path on the tree from the root, i.e., the Internet peering point. A storage cache can be located at each node of the network, providing a potential facility for storing data. Moreover, caches are organized in a hierarchical structure: if a requested content is not available in a given cache, the request is forwarded to the parent cache of the hierarchy, without any collaboration among the caches located in the same level of the tree. Finally, we do not impose a given cache size, i.e., the cache size is an output of our approach.¹

The content distribution procedure is divided in the following steps: a) the content is fetched from the peering point to the storage caches located at a given level of the tree, b) the content is cached for a fixed amount of time, c) during this period the content is retrieved by users, based on its popularity. We then associate an energy cost to each of these steps, and we compute the total energy consumption. Our aim is then to find the optimal amount of data to cache at each level of the tree in order to minimize the overall energy consumption.

Focusing on power requirements, we consider the cost of keeping the content stored in the cache, the cost of reading/writing the content from/to the cache and the cost of sending the content through one hop of the tree. We assume that the cost for traversing one hop is different for each level, due to the different switching devices deployed in each

segment of the network [13]. In order to model the power consumption of each device, we assume a linear dependency with traffic volume, following the assumptions of previous works [12], [14], [15]. In particular, the cost of transporting information is expressed in terms of energy per bit, i.e., the total power consumption divided by the average throughput.

More formally, the set of levels in the network is $\mathcal{L} = \{1, \dots, L\}$, $L = |\mathcal{L}|$ being the number of levels. The peering point is located at level 1, while the users are connected to level L (e.g., the DSLAMs). We denote the total number of switching devices located at level $j \in \mathcal{L}$ as N_D^j . Let us define the storage cost for a single cache as C_S . C_R is the cost of reading/writing content on one cache. C_H^j is the cost of traversing one node located at level j in the network. Moreover, we consider the characteristics of the content. We assume that the content is represented by videos watched by users. τ is the total throughput of videos requested by users. Let us denote the average video size as A and the popularity window duration as I . Thus, the total number of videos V_W watched during the popularity window is:

$$V_W = \frac{\tau I}{A} \quad (1)$$

Let us define V_S as the total number of videos provided by the CP. We divide the videos into classes according to their popularity, N_C being the number of classes. The set of classes is denoted as $\mathcal{K} = \{1, \dots, N_C\}$. Class 1 is the most popular while class N_C is the least popular. We assume that, on average, each class has the same number of videos, which we denote as $V_C = \frac{V_S}{N_C}$.

For each class $k \in \mathcal{K}$ we adopt the Zipf-based popularity model of [16] and compute the number of videos watched per class as

$$V_W^k = V_W \frac{k^{-\beta}}{\sum_{k=1}^{N_C} k^{-\beta}} \quad (2)$$

β being the parameter of the Zipf distribution.

We then compute the energy consumed for disseminating class k when it is stored on the caches located at level j . In particular, we first compute the energy consumed for fetching the content into the caches, and to keep the content stored:

$$\phi(j) = AV_C N_D^j \left(\sum_{z=1}^{j-1} C_H^z + C_R + C_S I \right) \quad (3)$$

The first term inside parentheses is the cost of traversing $(j-1)$ hops. The second term is the cost of writing the content on the cache. The third term is the cost of keeping the content stored, which is multiplied by the popularity window duration I since this cost has to be always accounted for the whole time period. All the costs are then multiplied by the amount of information that it is stored in level j , i.e., $A \times V_C \times N_D^j$. Note that $\phi(j)$ does not depend directly on the popularity of the class but only on the level j chosen for caching.

We then compute the energy consumed for retrieving the content as:

$$\varphi(j, k) = AV_W^k \left(C_R + \sum_{z=j}^L C_H^z \right) \quad (4)$$

¹In our scenarios the obtained cache size is always lower than the maximum capacity of current storage devices.

In particular, we consider the cost of reading the content and the cost of sending the content from the caches at level j to users. The retrieved information corresponds to the videos that are watched during the popularity window duration, i.e., $A \times V_W^k$. Differently from $\phi(j)$, $\varphi(j, k)$ depends on both the class popularity and the level where the content is cached.

The total energy consumed for disseminating class k on level j is:

$$E_k^j = \begin{cases} \phi(j) + \varphi(j, k), & j > 0 \\ AV_W^k \sum_{z=1}^L C_H^z, & j = 0 \end{cases} \quad (5)$$

Note that level 0 is the special case where the data is served from the original source, i.e., caching is not exploited within the considered network. In this case, the total energy consumption is the cost of sending the watched videos directly from the peering point to the users.

The best level to store the videos of class k is then:

$$h_k = \underset{j \in \mathcal{L}}{\operatorname{argmin}} E_k^j \quad (6)$$

Note that the best level for each class is computed independently from the other classes. We therefore repeat this procedure for each class k .

The total energy consumption with caching is computed as:

$$T = \sum_{k \in \mathcal{K}} E_k^{h_k} \quad (7)$$

Which we can compare to the energy consumption without caching:

$$T' = \sum_{k \in \mathcal{K}} E_k^0 \quad (8)$$

By comparing T with T' , we can estimate if caching is effective or not in saving energy. However, computing Eq.(6) for each class is not feasible, since the iteration over the levels has to be repeated for all the classes, resulting in a time complexity of $\mathcal{O}(L \times N_C)$. To solve this issue, we have proposed a new algorithm in order to efficiently compute T .

III. GCT ALGORITHM DESCRIPTION

We first detail the properties that we have exploited to design our algorithm. In particular, since there is no limit on the storage, we can choose the best level for every video class independently from the others. Moreover, for level j that is optimal for the video class k , we have necessarily $E_k^j \leq E_k^{j'}$ for any j' different from j . This implies:

$$\varphi(j, k) - \varphi(j', k) \leq \phi(j') - \phi(j). \quad (9)$$

In addition to that, when two video classes are stored in the same level, less videos will be retrieved from this level for the less popular class. Namely, if a class k' is less popular than another class k , then: $\varphi(j, k') \leq \varphi(j, k)$ and $\varphi(j', k') \leq \varphi(j', k)$. This leads to the following inequality:

$$\varphi(j, k') - \varphi(j', k') \leq \varphi(j, k) - \varphi(j', k). \quad (10)$$

The two previous equations imply:

$$\varphi(j, k') - \varphi(j', k') \leq \phi(j') - \phi(j). \quad (11)$$

And hence we get the following property.

Property 1. Let k and k' be two video classes such that k -class videos are more popular than k' -class ones ($k' > k$) and let j be the optimal level for k . Then for every level j' lower than j ($j' < j$) we have $E_{k'}^j \leq E_{k'}^{j'}$.

The intuition of the Green Content Threshold (GCT) algorithm is to restrict the evaluation of Eq. (6) to specific k , which we call thresholds. A threshold is defined as the last class to be stored at level j , before starting storing in another level x (with $x < j$). The rule for deciding when to pass from one level to another one is based on the energy consumption E_k^j (recall that level 0 correspond to the case without caching). In particular, we find the class index $\bar{k} = \bar{k}(j, x)$ that verifies, for arbitrary levels j and x the following equality:

$$E_{\bar{k}}^j = E_{\bar{k}}^x \quad (12)$$

For some classes, the caching of their videos in level j is preferred to caching them in level x in term of energy efficiency; and for some other classes it is the opposite. $\bar{k}(j, x)$ is the index that separates both set of classes. In fact, Eq. (12) being verified by \bar{k} implies:

$$\varphi(x, \bar{k}) - \varphi(j, \bar{k}) = \phi(j) - \phi(x). \quad (13)$$

On the other side $k \leq \bar{k}$ iff:

$$\varphi(x, \bar{k}) - \varphi(j, \bar{k}) \leq \varphi(x, k) - \varphi(j, k). \quad (14)$$

And this leads to the following property.

Property 2. Let x and j be two levels such that $x < j$ and let \bar{k} be the solution to $E_{\bar{k}}^j = E_{\bar{k}}^x$. Then $k \leq \bar{k}$ iff $E_k^x \geq E_k^j$.

To compute $\bar{k}(j, x)$ we solve Eq.(12) using Eq.(5), obtaining:

$$\bar{k}(j, x) = \begin{cases} \left[\frac{V_W (\sum_{z=x}^j C_H^z)}{(\Delta_f(j, x)) \sum_{k=1}^{N_C} k^{-\beta}} \right]^{\frac{1}{\beta}}, & j > 0 \\ \left[\frac{V_W (\sum_{z=x}^j C_H^z - C_R)}{f(j) \sum_{k=1}^{N_C} k^{-\beta}} \right]^{\frac{1}{\beta}}, & j = 0 \end{cases} \quad (15)$$

with $f(j) = N_D^j \left(\sum_{z=1}^j C_H^z + C_R + C_{SI} \right)$ and $\Delta_f(j, x) = f(x) - f(j)$.

The matrix with elements $\bar{k}(j, x)$ is denoted as \bar{K} . Each element of this matrix represents a threshold class for moving from one level to another one.

The algorithm that we propose is then divided into three steps: a) computation of the thresholds matrix, b) best threshold selection, and c) computation of total power consumption.

The first step is performed by computing \bar{K} with Eq.(15) for each $j \in \mathcal{L}$ and each $x < j$. In the next step, we select the best threshold for each level, by adopting the procedure reported in Alg.1. The function takes as input the matrix \bar{K} , the number of levels L , and the number of classes N_C . The array of best thresholds B is produced as output. The algorithm searches the best thresholds from the lower levels (i.e., the access nodes) to the upper ones. In particular, the minimum threshold is selected by evaluating $\bar{k}(j, x)$ from the current level to each upper level (lines 5-10): in fact, thanks to property

Algorithm 1 Pseudo-Code of the best threshold selection procedure.

Input: threshold matrix \overline{K} , number of levels L , number of classes N_C ;
Output: array of best thresholds B

```

1: curr_level = L;
2: B [0]=N_C- 1;
3: while curr_level != 1 do
4:   min_thre=inf;
5:   for upper_level = 1:curr_level-1 do
6:     curr_thre =  $\overline{K}$  [curr_level,upper_level];
7:     if curr_thre < min_thre then
8:       min_thre=curr_thre;
9:     end if
10:  end for
11:  B [curr_level]=min_thre;
12:  curr_level=curr_level-1;
13: end while

```

1, we can ignore levels that are lower than the current level. Moreover, we know, thanks to property 2, that for every video that has a popularity rank lower than the minimum threshold it is better to cache it in current level than in any higher level. The procedure is repeated until the last level is reached (line 3). It is clear that when the values of \overline{K} are computed optimally, this algorithm is optimal.

We then detail how the total energy consumption is computed from the best thresholds. We first derive the energy consumption consumed at level j . This term includes the energy consumption corresponding to the classes that are assigned to the current level j , i.e., $b_{j-1} < k \leq b_j$ (b_{j-1} and b_j being elements of the array of the best thresholds B), which can be expressed as:²

$$\sum_{k=(b_{j-1}+1)}^{b_j} E_j^k = [b_j - b_{j-1} + 1] \phi(j) + \sum_{k=(b_{j-1}+1)}^{b_j} \varphi(j, k) \quad (16)$$

In particular, from the definition of $\varphi(j, k)$ the last term can be expressed as:

$$\sum_{k=(b_{j-1}+1)}^{b_j} \varphi(j, k) = AV_W N_D^x \left(C_R + \sum_{z=j}^L C_H^z \right) \frac{\sum_{k=(b_{j-1}+1)}^{b_j} k^{-\beta}}{\sum_{k=1}^{N_C} k^{-\beta}} \quad (17)$$

The exponential terms can be expressed as:

$$\sum_{k=a}^c k^{-\beta} = \zeta(\beta, a) - \zeta(\beta, c+1) \quad (18)$$

ζ being the Hurwitz zeta function [17]. To compute the total energy consumption T , the algorithm solves Eq.(16)-(18) for all the levels, and the sum of the energy consumption is stored in T .

GCT has a time complexity that depends on the time needed to evaluate $\sum_{k=a}^c k^{-\beta}$ for $0 \leq a \leq c \leq N_C$. When this time is bounded by \mathcal{T} , the complexity is $\mathcal{O}(L^2 + L \times \mathcal{T})$. Since there are efficient ways to approximate ζ , we can have a good

²A similar expression can be derived for the $j = 0$ case (with $b_{-1} = 0$).

approximation of $\sum_{k=a}^c k^{-\beta}$ and of GCT with a low time complexity. In particular, if there exists a ρ -approximation of ζ which has a time complexity \mathcal{T}_ρ , then we can get a ρ -approximation on the optimal solution (the minimum energy consumption induced by all classes) that have a time complexity $\mathcal{O}(L^2 + L \times \mathcal{T}_\rho)$ which is better than the original approach since normally $L \ll N_C$ and $\mathcal{T}_\rho \ll N_C$.

IV. RESULTS

We have implemented the GCT algorithm in Python. In particular, we have adopted the `mpmath` library for an efficient computation of the Hurwitz Zeta function ζ .³ We have then evaluated GCT over two realistic networks of national ISPs, namely France Telecom (FT) and an ISP in Morocco. The main features of the networks, together with the setting of the main parameters, are reported in Tab. I. Both networks are composed of six levels in total (core, core-regional, metro-core, metro, access-metro, access), and with a different number of switching devices deployed at each level.

Focusing on power requirements, the cost of storage C_S is taken from [18]. The cost of reading/writing the cache is provided privately by CloudFlare Inc [19], based on their global network of content caches. The costs of network hops are based on equipment datasheets, measurements and operational conditions published in [20] and provided by France Telecom.

We then consider the characteristics of the content. In particular, we assume that for the FT scenario traffic forecasts are provided for the year 2020. We refer the reader to [21] for a detailed description on how these forecasts are obtained. On the contrary, for the Moroccan scenario we set values in accordance to the current traffic measured over the network. As a consequence, the total video throughput τ is eight times higher in the FT scenario compared to the Moroccan one. To this extent, we have also considered different values for the average video size A , assuming for the FT scenario a value that corresponds to a high definition video provided today on optical disks. Moreover, we have assumed a popularity duration of one week,⁴ and a value for the exponent of the Zipf distribution from [16] for both scenarios. Finally, we assume a number of videos of a typical video CP, and one video for each class.

Evaluation Metrics We describe the metrics adopted to evaluate the performance of our algorithm. We first derive the energy saving as:

$$S = \frac{T' - T}{T'} \quad (19)$$

We then compute the percentage of bandwidth that is saved at the peering point when caches are exploited:

³The complexity of the implemented Hurwitz function is in the order of $\mathcal{O}(p^{2+\epsilon})$, p being the precision (the number of significant bits) and ϵ a small number.

⁴Even though the popularity of the watched videos can actually change during this time period, the popularity of the most viewed videos is almost constant. As an example, [16] shows that the popularity of the 50 most viewed videos does not consistently vary over the days. In this work, we are interested in the most popular videos, as these contents are cached inside the ISP network.

TABLE I
MAIN PARAMETERS FOR THE CONSIDERED SCENARIOS.

Parameter	FT	Moroccan
L [units]	6	6
N_D^j [units]	[1 8 24 216 216 2160]	[1 20 20 20 200 10000]
C_S [W/Gb]	$9.375 \cdot 10^{-4}$	$9.375 \cdot 10^{-4}$
C_R [J/Gb]	24.3	24.3
C_H^j [J/Gb]	[12.5 25 30 35 200 300]	[12.5 25 30 35 200 300]
τ [Gb/s]	$8 \cdot 10^3$	10^3
I [days]	7	7
A [Gb]	15	0.6
β [units]	0.8	0.8
V_S [units]	$120 \cdot 10^6$	$120 \cdot 10^6$
N_C [units]	V_S	V_S

TABLE II
SUMMARY OF RESULTS FOR THE TWO NETWORK SCENARIOS

Metric	FT	Moroccan
Energy savings (S)	8.7%	11.0%
Yearly monetary savings [k€]	769	122
Bandwidth savings (Ψ)	18.2%	30.2%
Cache Size [GB]	Λ^1	0
	Λ^2	0
	Λ^3	32546
	Λ^4	0
	Λ^5	23510
	Λ^6	35878
		5581
		2041
		46

$$\Psi = \frac{\tau I - \sum_{k:\{h_k=0\}} AV_W^k - \sum_{k:\{h_k>0\}} AV_C N_D^{h_k}}{\tau I} \quad (20)$$

In particular, τI is the total amount of information flowing through the peering point without caching; $\sum_{k:\{h_k=0\}} AV_W^k$ is the amount of watched videos not stored inside the network in the case with caching; $\sum_{k:\{h_k>0\}} AV_C N_D^{h_k}$ is the total amount of information initially fetched inside the caches.

Finally, we consider the cache size for a device in level j , defined as:

$$\Lambda^j = \sum_{k:\{h_k=j\}} AV_C \quad (21)$$

General Analysis Tab. II reports the results for the two scenarios obtained with the GCT algorithm. We first consider the energy savings compared to the case in which caching is not exploited. Energy savings of almost 9% and 11% are possible for the FT and the Moroccan scenarios, respectively. By assuming that caches are refreshed once a week for an entire year, we have estimated monetary savings⁵ of more than 700 k€ for FT, and more than 100 k€ for the Moroccan network. Moreover, the savings in terms of bandwidth saved at the peering point are even larger, reaching 18% for the FT scenario and 30% for the Moroccan one.

The table also reports the cache size Λ^j per device for each level j . Interestingly, Λ^j is at most 36 TB, a value that can be covered by a commercial array of disk drives. Moreover, the capacity requirements tend to decrease moving closer to users, with at most 2041 GB of storage required at the access

⁵We have assumed an electricity cost of 0.21€/kWh.

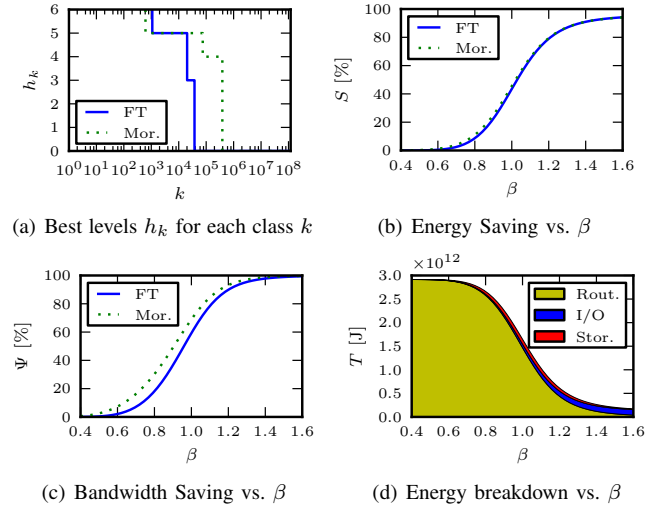


Fig. 1. Best levels for each class k (a) and impact of the variation of β (b-d).

level for the FT network and only 46 GB for the Moroccan one.

To give more insight, Fig. 1(a) reports the best level h_k for each class k for the two scenarios. k ranges between 1 and V_C . The levels on the left are the most popular ones and hence, to minimize the costs of moving the information frequently from the cache to users, it is better to store these classes in the closest level to users, i.e., the access part of the network. Moving from left to right, the popularity decreases, and therefore the classes are stored in the inner levels of the topology (metro and core). At last, very unpopular classes are assigned to level 0, i.e., they are not cached at all. Interestingly, the percentage of the total number of stored classes is around 1.7% and 0.5% for the Moroccan and the FT networks, respectively. Thus, we can conclude that with the considered power and popularity models, the ISP needs to store a little amount of content information to achieve energy and bandwidth savings. This is an encouraging result showing that caching not only has benefits on QoS and customer experience, but it can also lead to a better management of the ISP power consumption.

Impact of Content Characteristics We then consider how much the characteristics of the content impact the energy and the bandwidth savings. We first vary one parameter per time, keeping the others with the default values reported in Tab. I.

We start considering the variation of β , as reported in Fig. 1(b) and Fig. 1(c). With low values of β (left part of the figures) all the classes tend to have a similar popularity. Intuitively, there is not a huge benefit in terms of energy in storing these classes inside the network, since the cost of storing this amount of information would be huge. On the contrary, when β takes higher values (right part of the figures), the variation on the popularity also increase. This means that few classes have a very high popularity, while most of the content is seldom accessed by users. This in turn imposes to store the most popular classes close to users, and therefore both the energy and the bandwidth savings steadily increase. At last, when $\beta = 1.6$, the bandwidth savings are almost equal


Fig. 2. Impact of the total number of videos watched V_W .

TABLE III
VARIANCE DECOMPOSITION OF S FOR THE CONTENT PARAMETERS (FT SCENARIO)

Parameter	First Order Index	Total Effect Index
β	0.9950	0.9988
τ	0.0004	0.0007
V_S	0.0001	0.0002
A	0.0000	0.0000

to 100%, while energy savings are more than 90%. This corresponds to the case in which the most popular contents are cached in the last level of the network (i.e., the access part), and the main cost incurred by the ISP is to transfer these contents from the caches to the users. Thus, we can conclude that β greatly influences the performance in the network.

To give more insight, Fig.1(d) reports the breakdown of energy for the FT scenario considering: a) the energy consumed to route the content inside the ISP network, either from the peering point to the cache or from the cache to users ('Rout.' label), b) the energy consumed for reading/writing the content from/to the caches ('I/O' label), c) the energy consumed for keeping the content stored ('Stor.' label). For $\beta = 0.4$ the caching is not exploited, and therefore the largest amount of energy is due to the routing, i.e., the cost of moving information inside the network. However, as β increases, the routing energy steadily decreases, since caches are more frequently used. This in turn implies that the energy consumed for reading and writing the content on the caches increases. However, the total energy consumption is always decreased, producing high energy savings.

In the following, we consider the variation of the total number of videos watched V_W , reported in Fig. 2. When V_W increases (right part of the figures), both the energy and the bandwidth saving increase. This is due to the fact that as V_W increases the caches are more frequently accessed by users, and therefore the introduced gain in terms of energy and bandwidth is higher. Clearly, when the users seldom access the content (left part of the figures) there is no need to put caches.

Finally, we have applied a variance-based sensitivity analysis [22] to precisely characterize the relative impact of the parameters. In particular, we have considered how much the variance of the energy saving S is impacted by the variation of the content parameters. To this end, we have considered the first order index and the total effect index. The first order index takes into account how much the variance of a single parameter influences the variance of the output. On the contrary, the total effect index takes into account the effects of varying the parameter on the model's output, including all the variances


Fig. 3. Impact of the cost for reading and writing C_R .

TABLE IV
VARIANCE DECOMPOSITION OF S FOR THE POWER CONSUMPTION PARAMETERS (FT SCENARIO)

	First Order Index	Total Effect Index
C_S	0.2738	0.2699
C_R	0.0303	0.0265
C_H^1	0.1942	0.2009
C_H^2	0.3062	0.3125
C_H^3	0.0367	0.0436
C_H^4	0.0152	0.0210
C_H^5	0.0066	0.0143
C_H^6	0.1251	0.1259

from interaction with the other parameters. To compute both indexes, we have adopted a Monte Carlo method. In particular, we have generated a pool of 18 million samples, in which each parameter take a random value in the interval $[1/2, 2]$ w.r.t. the standard values reported in Tab. I.

Tab. III reports the computed indexes considering the total energy savings. The largest contribution to the output variance of the first order index is due to β , while the other parameters play a minor role. Moreover, the average video size does not impact the energy savings since this parameter is simplified when computing S . The table also reports the values of the total effect index. These values are very similar to the ones of the first order index. This shows that simultaneous varying multiple input variables does not have a strongly amplified (multiplicative) effect on energy savings when compared to varying them separately.

Impact of Power Consumption Models We first consider the variation of the cost for reading and writing C_R . Fig. 3 reports the saving in terms of energy and bandwidth. As expected, the savings are increasing when C_R decreases. Intuitively, the lower is the cost for reading and writing information, the higher is the gain introduced by caching. In particular, when $C_R \approx 0$ energy savings of more than 9% and 12% are possible for the FT and the Moroccan networks, respectively. Thus, we can expect that, if the energy efficiency of caches improves faster than the one of transport equipment, the benefit introduced by caching will be greater in the future.

Finally, we have performed the variance decomposition analysis also for the power consumption parameters, adopting the same procedure as in the previous subsection. Tab. IV summarizes the main results for the FT scenario. Interestingly, the energy savings are greatly impacted by the cost of storing the content C_S . This is due to the fact this term has to be counted for all the time periods and for all the caches, thus at the end its contribution is not negligible. Thus, it is very important to deploy energy efficient storage inside the ISP to




Fig. 4. Impact of the number of levels and the average degree.

obtain energy savings. Moreover, the energy costs in the first levels are also impacting the energy savings, since most of traffic reduction occurs in these levels.

Impact of Network Properties To give more insight, we have considered a network with the same degree for all levels, and we have studied the impact of the variation of the degree and the number of levels. In this way, we are able to study the impact of caching over a set of topologies. In particular, the degree of level j is defined as the average number of links connecting a device in level j with the devices in level $(j + 1)$. Moreover, we assume a number of video requests proportional to the degree and the number of levels. Specifically, we have assumed 5000 users connected for each device in the last level (i.e., the access one)⁶, each of them watching 3 videos of size $A = 15$ Gb during a popularity window duration I of one week. Additionally, we have set $V_S = 120 \times 10^6$. Focusing on power requirements, we have set $C_S = 9.375 \times 10^{-4}$ W/Gb, $C_R = 24.3$ J/Gb, $C_H = 25$ J/Gb for each level j , respectively. Fig. 4 reports the energy and the bandwidth savings. Interestingly, the savings increase with the number of levels and the average degree. In particular, energy savings of more than 20% are possible for very large networks. Moreover, bandwidth savings quickly approach 100%. The fact that energy and bandwidth savings increase with the degree and the number of levels L is due to two main reasons: i) increased cost of moving information inside the network when the number of levels is increased, and ii) total number of watched video increased, while total number of stored videos kept constant.

V. CONCLUSIONS

We have studied the energy-efficient design of a content architecture in a ISP network, by exploiting caches managed by the ISP. Our results indicate that caching brings substantial savings in terms of energy and bandwidth.

As next steps, we will consider the joint management of the content distribution architecture. In particular, our aim is to study the traffic variation over time and to compute the best set of caches powered on to satisfy a given traffic demand, while leaving the remaining caches powered off. Another possible direction is to introduce cooperation between neighboring caches to serve users and to reduce the amount of stored information. Finally, we plan to study the impact of

considering more than one peering node and the impact of introducing realistic traffic matrices inside the ISP.

REFERENCES

- [1] S. Lambert, W. Van Heddeghem, W. Vereecken, B. Lannoo, D. Colle, and M. Pickavet, "Worldwide electricity consumption of communication networks," *Optics Express*, vol. 20, no. 26, pp. 513–524, 2012.
- [2] C. Lange, D. Kosiankowski, R. Weidmann, and A. Gladisch, "Energy consumption of telecommunication networks and related improvement options," *Selected Topics in Quantum Electronics, IEEE Journal of*, vol. 17, no. 2, pp. 285–295, 2011.
- [3] "Cisco visual networking index: Forecast and methodology, 2011–2016," 2012.
- [4] R. Bolla, R. Bruschi, F. Davoli, and F. Cucchietti, "Energy efficiency in the future internet: a survey of existing approaches and trends in energy-aware fixed network infrastructures," *Communications Surveys & Tutorials, IEEE*, vol. 13, no. 2, pp. 223–244, 2011.
- [5] L. Chiaraviglio and I. Matta, "Greencoop: cooperative green routing with energy-efficient servers," in *Proceedings of the 1st ACM International Conference on Energy-Efficient Computing and Networking*, Passau, Germany, 2010, pp. 191–194.
- [6] U. Lee, I. Rimac, and V. Hilt, "Greening the internet with content-centric networking," in *Proceedings of the 1st ACM International Conference on Energy-Efficient Computing and Networking*, Passau, Germany, 2010, pp. 179–182.
- [7] V. Valancius, N. Laoutaris, L. Massoulié, C. Diot, and P. Rodriguez, "Greening the internet with nano data centers," in *ACM CoNEXT '09*, Rome, Italy, 2009, pp. 37–48.
- [8] J. Llorca, K. Guan, G. Atkinson, and D. Kilper, "Energy efficient delivery of immersive video centric services," in *IEEE INFOCOM*, Orlando, USA, 2012, pp. 1656–1664.
- [9] P. Gill, M. Arlitt, Z. Li, and A. Mahanti, "The flattening internet topology: Natural evolution, unsightly barnacles or contrived collapse?" in *Passive and Active Network Measurement*. Cleveland, USA: Springer, 2008, pp. 1–10.
- [10] L. Sofman and B. Krogfoss, "Analytical model for hierarchical cache optimization in iptv network," *Broadcasting, IEEE Transactions on*, vol. 55, no. 1, pp. 62–70, 2009.
- [11] C. Jayasundara, A. Nirmalathas, E. Wong, and C. Chan, "Energy efficient content distribution for VoD services," in *Optical Fiber Communication Conference*, Los Angeles, USA, 2011, pp. 1–3.
- [12] U. Mandal, C. Lange, A. Gladisch, P. Chowdhury, and B. Mukherjee, "Energy-efficient content distribution over telecom network infrastructure," in *13th International Conference on Transparent Optical Networks (ICTON)*, Stockholm, Sweden, 2011, pp. 1–4.
- [13] U. Lee, I. Rimac, D. Kilper, and V. Hilt, "Toward energy-efficient content dissemination," *Network, IEEE*, vol. 25, no. 2, pp. 14–19, 2011.
- [14] J. Baliga, R. Ayre, K. Hinton, and R. Tucker, "Architectures for energy-efficient iptv networks," in *Optical Fiber Communication Conference*, San Diego, California, 2009.
- [15] A. Feldmann, A. Gladisch, M. Kind, C. Lange, G. Smaragdakis, and F.-J. Westphal, "Energy trade-offs among content delivery architectures," in *IEEE CTTE '10*, 2010, pp. 1–6.
- [16] M. Cha, P. Rodriguez, J. Crowcroft, S. Moon, and X. Amatriain, "Watching television over an IP network," in *ACM IMC '08*, Vouliagmeni, Greece, 2008, pp. 71–84.
- [17] A. Voros, "Zeta functions for the riemann zeros," in *Annales de l'institut Fourier*, vol. 53, no. 3, 2003, pp. 665–700.
- [18] OCZ Technology Group. [Online]. Available: <http://www.ocztechnology.com/ocz-revdrive-3-x2-pci-express-ssd.html>
- [19] Cloud Flare Inc. [Online]. Available: <https://www.cloudflare.com/>
- [20] W. Van Heddeghem, F. Idzikowski, E. Le Rouzic, J. Mazeas, H. Poignant, S. Salaun, B. Lannoo, and D. Colle, "Evaluation of power rating of core network equipment in practical deployments," in *IEEE Online conference on green communications (GreenCom)*, 2012.
- [21] O. Renais and J. Le Briand, "Tracks for transport network architecture optimization," in *NETWORKS 2012*, Rome, Italy, 2012, pp. 1–6.
- [22] A. Saltelli, M. Ratto, T. Andres, F. Campolongo, J. Cariboni, D. Gatelli, M. Saisana, and S. Tarantola, *Global sensitivity analysis: the primer*. Wiley-Interscience, 2008.

⁶Note that while the number of user per device is constant, the total number of users scales with the degree and the number of levels L .