

Optimizing Server Power Consumption in Cross-Domain Content Distribution Infrastructures

Chang Ge, Ning Wang, Zhili Sun
Centre for Communication Systems Research
University of Surrey, Guildford, UK
{C.Ge, N.Wang, Z.Sun}@surrey.ac.uk

Abstract — Optimizing server’s power consumption in content distribution infrastructure has attracted increasing research efforts. The technical challenge is the tradeoff between server power consumption and the content service capability on both the server and the network side. This paper proposes and evaluates a novel approach that optimizes content servers’ power consumptions in large-scale content distribution platforms across multiple ISP domains. Specifically, our approach strategically puts servers to sleep mode without violating load capacities of virtual content delivery links and active servers in the infrastructure. Such a problem can be formulated into a nonlinear programming model. The efficiency of our approach is evaluated in a content distribution topology covering two real interconnected domains. The simulation has shown that our approach is capable of reducing servers’ power consumptions by up to 62.2%, while maintaining the actual service performance in an acceptable scope.

I. INTRODUCTION

In a typical content distribution infrastructure covering one or more ISP domains, several content servers are strategically established to serve incoming content requests, *e.g.*, in a localized manner. As illustrated in Fig. 1, content servers are often planted close to some Point-of-Presence (PoP) nodes with high content demand volumes to reduce bandwidth cost in content delivery [1]. Normally, content requests are resolved to either local (preferred) or remote servers, and the contents are effectively delivered through the virtual path from the server PoP and the user PoP. The determination of the virtual paths is based on the underlying intra- and inter-domain routing configurations. In order to support end-to-end QoS assurance in content delivery, the content provider may “rent” bandwidth resources along the virtual links in the overlay across PoP nodes through the establishment of Service Level Agreements between content providers and underlying ISPs. Such bandwidth reservation by content providers from the underlying ISP networks can be typically achieved via virtualization techniques where “virtual” bandwidth capacity can be reserved by content providers.

Energy efficiency in the Information Communication Technologies (ICT) sector has been more and more emphasized in recent years, which aims at the development of future green network infrastructures *e.g.*, Nano Data Centers [2]. More recently, authors in [3] have considered optimizing both content providers’ and ISPs’ energy costs simultaneously assuming shared information between them. Later, a fully

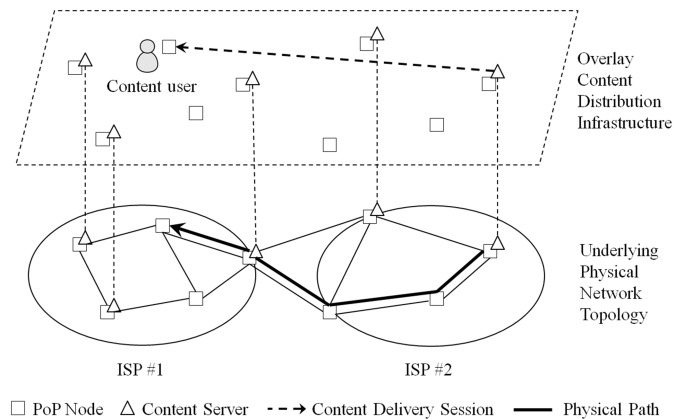


Figure 1. Overlay content distribution infrastructure

distributed approach was proposed in [4] without such an assumption by employing dual decomposition technique. Near-optimal energy saving performance was achieved, but at the cost of considerably high computational complexity.

As far as content providers are concerned, it can be desired to put some servers to sleep when incoming content request demands are low, for instance during off-peak hours. However since content servers are geographically distributed in content delivery infrastructures, how to efficiently save their power consumption is complicated, as the decision on servers’ sleep mode reconfiguration should not only take into account the service capability *at the server side*, but also *network resource availability* to support end-to-end content delivery.

In this paper, we propose a novel scheme that aims to optimize content servers’ power consumptions in cross-domain content distribution infrastructure by strategically putting servers to sleep without violating service constraints on both the server side and the network side. Specifically, based on the captured content consumption request pattern during off-peak time [5], a subset of content servers attached to some PoP nodes can be reconfigured to sleep mode. In such a scenario, ongoing active sessions need to be diverted to the remaining servers that are still awake, but located in different sites. On one hand, the optimization scheme tries to put as many servers to sleep as possible provided that the overall remaining server capacities are able to handle the ongoing sessions. On the other hand, content server switching also leads to changes in the bandwidth utilization of the underlying network, as some consumers need to fetch the content from

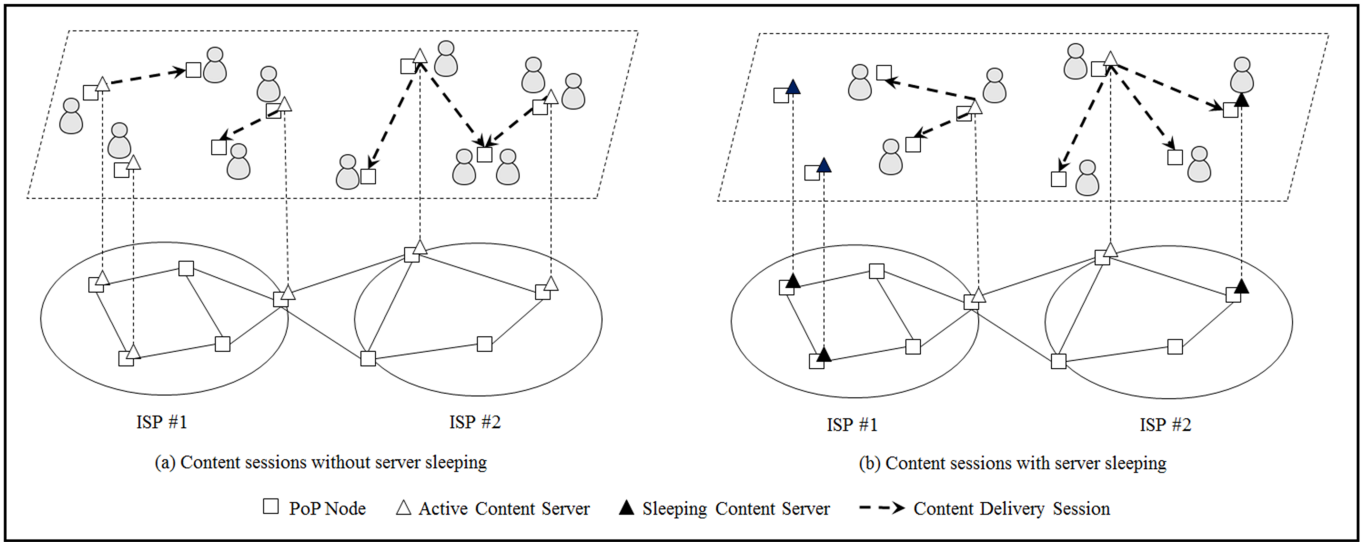


Figure 2. Working mechanism of the proposed scheme

more remote servers. Intuitively, the availability of fewer active servers will incur more bandwidth consumption by content traffic that is generated by redirecting content requests to remote servers. In such a scenario, specific constraints need to be taken into account which should keep both the server load and the traffic load on the virtual links under their respective capacities in order to avoid QoS deterioration.

We formulate the problem into a nonlinear programming model with the objective of minimizing content servers' power consumptions subject to both server capacities and network bandwidth availability constraints. It is worth mentioning that, such a scheme can be *independently* used by content providers for energy efficiency, depending on their own resources (servers and availability of rented bandwidth), but without necessary coordination with the underlying ISPs. This is in contrast to some existing approaches which need complex interaction (e.g., necessary context information sharing) between content providers and ISPs for joint optimizations [3]. To evaluate the power saving effect of our model and its impact on network performance, we used the realistic interconnected GEANT-Abilene inter-domain topology at PoP level [6] [7]. We mainly compare our results against the scheme in [8], which aims to minimize network bandwidth cost without considering content servers' power consumptions. We observed that our approach is capable of reducing servers' power consumptions by up to 62.2% without significantly affecting the overall service assurance and network performance.

II. APPROACH OVERVIEW

We consider widely distributed content distribution infrastructures that cover multiple autonomous ISP domains. Each domain contains several Point-of-Presence (PoP) nodes where local users send out requests for consuming content objects. Each incoming content request is assigned to distinct servers with specific policies, for instance locality [1]. Such a process is typically done via an overlay control plane maintained by the content provider.

Fig. 2(a) illustrates such an operation, in which five servers handle requests from nine PoP nodes within two domains. We can see that since the servers are attached to five of the PoP nodes, requests from these nodes can be handled locally without incurring inter-PoP delivery bandwidth cost. Requests from the other four nodes are directed to servers closest to them, which is often the case in practice to minimize user-experienced latency [1]. In order to offer end-to-end Quality of Service (QoS) support for content delivery to consumers, the content provider needs to "rent" bandwidth resources from the underlying involved ISP networks. Specifically, the overlay content distribution infrastructure is established based on the underlying PoP level topology, and the virtual bandwidth capabilities of the physical links are leased to the virtual links in overlay content distribution infrastructure. For simplicity, in this paper we assume the overlay topology has the same layout as the (inter-domain) PoP topology, however the proposed techniques can also be used for any generic content overlay topology. The routing of the content delivery paths in this overlay is effectively determined by the intra- and inter-domain routing configurations.

During normal period, all content servers are active for handling content requests, as shown in Fig. 2(a) where content requests are either served locally or by content servers nearby. However, keeping all content servers running may lead to considerable waste of power consumptions during off-peak hours. Firstly, the intensity of incoming content requests at a site varies across different times within a day, which becomes very low at midnight and early morning [5]. Secondly, since time-zone differences exist between PoP nodes (cities) which are distributed globally, it is unlikely to have high content request volumes at all nodes simultaneously. Based on these observations, it can be inferred that a subset of the content servers can be sufficient for handling all content requests during off-peak time, which may lead to considerable amount of power saving by putting the remaining servers to sleep mode. This is the main idea we propose in this paper.

Our approach on saving content servers' power is illustrated in Fig. 2(b). Assume that during off-peak time at some regions, the overall incoming content requests only require two active servers to handle them. Under such scenario, our approach takes as input the content distribution overlay topology (*i.e.*, nodes, links and server locations), user content request volume from each PoP node, and determines which server(s) can be put to sleep without violating: (1) the load capacities of the active servers, and (2) bandwidth constraint on the links involved in the content delivery paths between PoP nodes. As shown in the figure, the content provider now puts three of the servers to sleep mode according to the off-peak content request intensity. Requests that can be originally handled by these three servers are now redirected to the remaining two active servers, and three servers' power consumptions can be eliminated except very low sleeping power. It is worth mentioning that, redirecting requests from local content servers belonging to the same PoP to remote ones inevitably increase bandwidth costs on the inter-PoP links, and it is easy to infer that the decision on server sleeping and content request mapping should also take into account bandwidth availability on overlay links.

III. PROBLEM FORMULATION

We model each autonomous domain as a uni-directional graph $G = (V, E)$ where each vertex represents a PoP node and each edge represents an intra-domain link between two PoP nodes. Each intra-domain link in E is assigned with a specific IGP link weight for determining intra-domain paths. Such path subsequently determines the routing of the content overlay. Multiple domains are interconnected through dedicated inter-domain links, and typically an inter-domain link connects two PoP nodes belonging to different domains. For each intra- and inter-domain link that is mapped onto the overlay content distribution infrastructure, a proportion of the link capacity is leased to the overlay, and such virtual capacity for supporting the content delivery is denoted as U_l for overlay link l . Now we formally specify the nonlinear programming model that aims to minimize servers' power consumptions. The notations used in the model are listed in Table I.

$$\text{Minimize: } \sum_{s \in S} P_s \quad (1)$$

subject to:

$$P_s = f(\delta_s, w_s) = \begin{cases} P_{idle} + \alpha w_s, & \delta_s = 1 \\ P_{sleep}, & \delta_s = 0 \end{cases} \quad (2)$$

$$w_s \leq W_s \quad \forall s \in S \quad (3)$$

$$w_s = \sum_{v \in V} R_v^s \quad \forall s \in S \quad (4)$$

$$u_l \leq U_l \quad \forall l \in E \quad (5)$$

$$u_l = \sum_{v \in V} \sum_{s \in S} \delta_l^{vs} R_v^s \quad \forall l \in E \quad (6)$$

$$\sum_{s \in S} R_v^s = R_v \quad \forall v \in V \quad (7)$$

$$R_v = \sum_{i=1}^{N_v} R_v(i) \quad \forall v \in V \quad (8)$$

$$\delta_s \in \{0, 1\} \quad \forall s \in S \quad (9)$$

$$\delta_l^{vs} \in \{0, 1\} \quad \forall l \in E, \forall v \in V, \forall s \in S \quad (10)$$

The control variable in the model is δ_s , which is a binary parameter indicating whether server s is active or put to sleep in the output. The objective cost function is in equation (1), which minimizes overall server power consumptions across all servers. In equation (2), we employ the power consumption model for servers in [3], which considers modern servers that adapt their power with server. The model also assumes a server's idle power consumption to be 60% of its peak one. In contrast, servers in sleep modes have their typical power consumptions as low as 5 Watts [9]. These facts establish a piecewise power consumption model for content servers as in equation (2). We set $P_{sleep} = 5$ Watts, $P_{idle} = 200$ Watts and $\alpha = 40$ according to [9] and [3] respectively.

The cost function in equation (1) sums up results of equation (2) that calculates power consumption of each server s . The piecewise function in equation (2) makes the cost function and the model nonlinear. Constraints (3) and (5) limit servers and virtual links from exceeding their load capacities, which are calculated via (4) and (6) respectively. Constraint (4) calculates load on server s by aggregating all content volume that are mapped onto it. Constraint (6) calculates load

TABLE I. NOTATIONS IN THE MODEL

| | |
|-----------------|--|
| S | set of content servers |
| V | set of PoP nodes |
| E, l | set of links, $l \in E$ |
| w_s | load on server s |
| W_s | capacity of server s |
| $L(v, s)$ | path from v to s |
| δ_l^{ps} | binary indicating whether $l \in L(v, s)$ |
| u_l | content link load on l |
| U_l | virtual bandwidth capacity of l for content delivery |
| R_v | overall content traffic towards v |
| N_v | overall number of content requests at v |
| $R_v(i)$ | content traffic incurred by content request i at PoP node v |
| R_v^s | overall content traffic originated from s and destined to users in v |
| δ_s | binary indicating whether server s is on/off |
| P_s | power consumption of server s |
| P_{idle} | power consumption of server s when it is on and idle |
| P_{sleep} | power consumption of server s when it is in sleep mode |
| α | constant coefficient as in [3] |

on virtual link l by adding up all mapped content traffic demands whose path involves it. Constraint (7) indicates that user content requests from PoP node v can be mapped onto one or multiple servers. Constraint (8) shows that at PoP node v , the overall content traffic is calculated by aggregating bandwidth cost incurred by each content request. Constraint (9) shows the binary variable determining whether server s is to be active or sleeping with the values of 1 or 0. Constraint (10) shows the binary indicating whether link l belongs to the path from node v to server s i.e., $L(v, s)$, which is calculated by Dijkstra's Algorithm.

In order to solve the formulated nonlinear programming problem, we firstly calculated δ_l^{ps} based on the input of the network topology including nodes' and servers' locations and interconnecting links. We then implemented the model with LINGO software which is a powerful nonlinear model solver [11], and applied linearization technique and branch-and-bound method to find global optimal solutions of the model.

IV. PERFORMANCE EVALUATION

A. Experimental Setup

As previously mentioned, our scheme takes as input the network topology including PoP node locations and physical links interconnecting them, content distribution overlay platform with content server locations, and user content request volumes at each PoP node.

Firstly for the underlying physical network, we consider two interconnected domain topologies i.e., GEANT [6] (in the EU region) and Abilene [7] (in USA) networks, which have their PoP nodes widely distributed over Europe (22 nodes) and USA (11 nodes) respectively. Altogether there are 53 network links with various bandwidth capacities that interconnect them, including three inter-domain links.

Secondly for the content overlay platform, we set nine content servers that handle content requests, among which four are attached to Abilene and five to GEANT. This number is assumed based on the number of PoP nodes in each domain and the statement in [8] showing that four or five servers are near-optimal in terms of content delivery cost in networks with tens of nodes. For the servers' locations, since the content provider is unlikely to change its servers' locations in an established content distribution infrastructure [10], we consider as input a set of different server location scenarios. Firstly, we consider the scenario that content servers are attached to the PoP nodes with the highest expected content request volumes, as this is the default server placement strategy employed by many content providers to minimize delivery costs [1]. In other scenarios we randomly distribute nine servers in the topology, with four in Abilene and five in GEANT. Altogether ten server placement scenarios have been created for evaluations.

Thirdly, the number of content requests at each PoP node is based on not only the local population of the cities, but also time difference within a day. There are various proportions of users that are actively sending content requests at different hours, as shown in [5]. To best reflect the pattern of content requests across the PoP nodes, we selected three specific

TABLE II. CONTENT REQUEST VOLUME SCENARIOS

| # | Time within the Day | | % of Active Users in USA | % of Active Users in EU |
|----|-----------------------|--------------------|--------------------------|-------------------------|
| | USA [GMT-8, GMT-5] | EU [GMT, GMT+2] | | |
| #1 | [20:00, 23:00] | [04:00, 06:00] | 5% | 1% |
| #2 | [04:00, 07:00] | [12:00, 14:00] | 1% | 7.5% |
| #3 | [13:00, 16:00] | [21:00, 23:00] | 7.5% | 5% |

content request volume scenarios considering time differences between the PoP nodes as listed in Table II. In two of the selected scenarios, either of the two domains is in off-peak time. And in another scenario, the two domains both experiences relatively high content request volumes. We will refer to the three selected scenarios in short as “*volume scenarios*” in the context.

B. Simulation Procedure

In order to evaluate our proposed scheme's power saving effect and its impact on network performance in terms of content traffic distribution against available rented bandwidth resources, we carried out the simulations in three groups, each one corresponding to a request volume scenario in Table II. Hence each group has the same underlying network topology and the same content request volume at each node, but different content server locations.

As a reference, we also modified and implemented the integer linear programming model in [8] which was based on unicast delivery protocol. It aimed at minimizing delivery cost including overall network bandwidth and server bandwidth costs, and power saving was not considered by the authors. The modification includes performing linear relaxation on the original model by allowing content requests from a PoP node to be mapped onto one or multiple servers. Also instead of optimizing servers' numbers and locations, we regard them as given inputs in coherence with the ten server placement scenarios in our experimental setup. We will refer to the modified model in [8] as the “*reference scheme*” in context.

Simulations were carried out within our scheme and the reference scheme under the same scenarios for comparison. We choose the following metrics to evaluate.

- 1) *Power saving performance metrics*
 - Overall content servers' power consumptions.
 - Number of remaining active servers.
- 2) *Network performance metrics*
 - Average and maximum link load in percentage – “link” here refers to the virtual link established by content provider with leased bandwidth resources from ISPs' underlying networks.
 - Average and maximum distances travelled by content requests – “distance” here refers to aggregated IGP link weights along the paths, which are proportional to the actual delay according to the link weight settings of GEANT and Abilene.

In addition to the general averaged results from the three scenarios, we specifically have a closer look at the servers' load performances under the first server location scenario with different content request volumes, which is based on the default server placement strategy in practice as described in

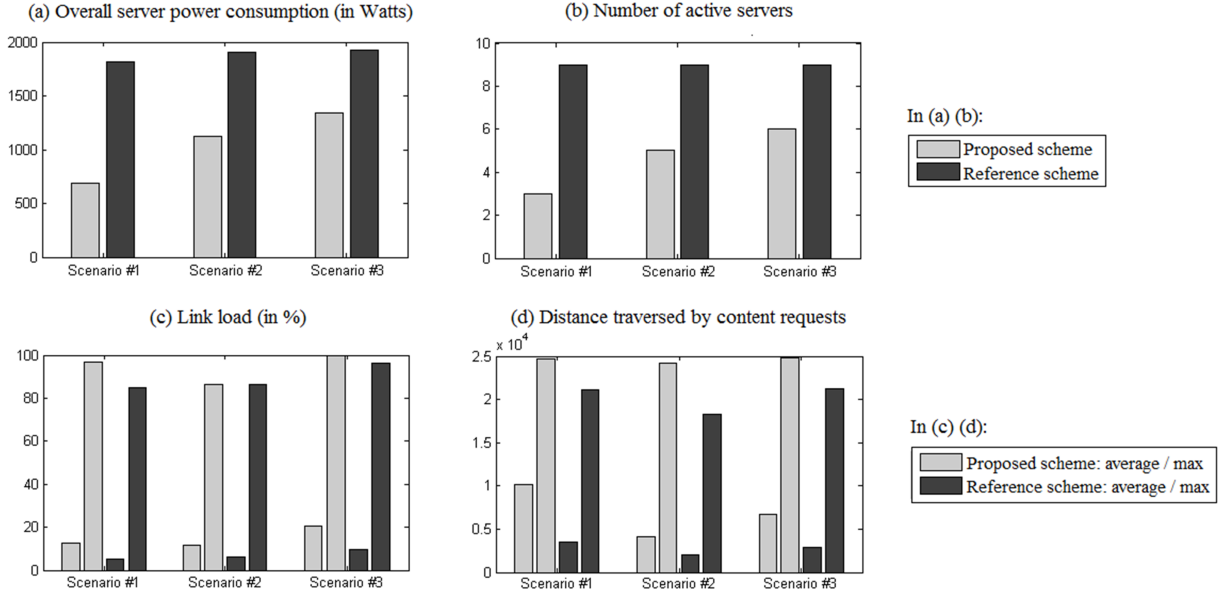


Figure 3. Simulation results

Section IV-A. Hence results under it would indicate our approach's performance in practice.

C. Results and Discussions

Figs. 3(a) and 3(b) indicate the overall power consumption performances. Figs. 3(c) and 3(d) indicate the network performance in terms of link load and network distance. All these performance metrics are compared against the scheme proposed in the reference scheme [8].

Fig. 3(a) shows that compared to the reference scheme, our approach is able to reduce the servers' power consumptions by 62.2%, 40.9% and 30.3% under volume scenarios #1, #2 and #3 respectively. Correspondingly Fig. 3(b) shows that respectively six, four and three servers can be put to sleep mode under each of the scenarios. With such reconfiguration, the ongoing content sessions can still have the service assurance with the remaining server capacity and the bandwidth support from the content distribution infrastructure with leased bandwidth.

Figs. 3(a) and 3(b) suggest that compared to the reference scheme that aims at minimizing delivery bandwidth cost without server power awareness, our scheme is able to achieve considerable conservation in power consumptions. From the simulation results, we observed that our approach is always capable of intelligently mapping content requests onto the remaining active servers under various server locations and content request volumes, meanwhile putting maximum number of servers to sleep for power saving. In contrast, the reference scheme keeps all servers active to minimize bandwidth costs by mapping more content requests to local servers. The results show that, when user content request volume is relatively low, there is large space for putting a subset of content servers to sleep mode while leaving the remaining ones to handle the overall content requests without impacting service assurance in terms of both content server capacity and virtual bandwidth support.

To examine the results in more detail, we look into individual server's load under the first server location scenario *i.e.*, the default scenario, in which servers are attached to PoP nodes with the highest local content request volumes. The actual load performance of individual server's load is shown in Table III. It can be observed that our scheme maps the incoming content requests to the least number of servers, while pushing the active servers' loads toward their capacities, but still able to handle all the incoming requests. In contrast, the reference scheme resolves the content requests in a relatively distributed manner to minimize bandwidth costs, but obviously all the content servers are kept active.

It can also be observed that in each volume scenario, our scheme tends to put specific servers with *low local content request volume* to sleep mode. The benefit is that the overall number of diverted content requests due to server sleeping can be reduced. In volume scenario #1, there is very low content request volume from the EU region and relatively higher demand in USA. Correspondingly the results show that four out of five content servers in GEANT are put to sleep mode. Similarly, in volume scenario #2 with low content demand volume in USA and relatively higher volume in the EU region, all servers in Abilene are put to sleep mode, and content requests from PoP nodes in USA are redirected to content servers in GEANT, but without violating the bandwidth availability constraint in the content overlay. Therefore, it can be inferred that while making decisions on servers sleep mode reconfiguration, our scheme takes servers' local content traffic volume into consideration such that PoP nodes with high traffic load are not affected by redirecting content requests to remote servers.

In addition to the power saving performance, we also evaluate the impact on network performance due to additional bandwidth costs (in terms of link load) incurred by redirecting content requests to remote servers. The metrics of virtual link loads and network distances traversed by content requests (*i.e.*,

TABLE III. SERVER LOAD PERFORMANCE UNDER DEFAULT SERVER PLACEMENT

| Server Domains and IDs | | GEANT (EU) | | | | | Abilene (USA) | | | |
|------------------------|-------------------------|------------|----------|----------|----------|----------|---------------|----------|----------|----------|
| | | Server 1 | Server 2 | Server 3 | Server 4 | Server 5 | Server 6 | Server 7 | Server 8 | Server 9 |
| Volume Scenario #1 | <i>Proposed Scheme</i> | 89.93% | Sleep | Sleep | Sleep | Sleep | 100% | Sleep | 39.33% | Sleep |
| | <i>Reference Scheme</i> | 34.17% | 5.5% | 19.83% | 1.67% | 13.67% | 11.67% | 31.67% | 84.67% | 26.33% |
| Volume Scenario #2 | <i>Proposed Scheme</i> | 100% | 44.17% | 100% | 100% | 100% | Sleep | Sleep | Sleep | Sleep |
| | <i>Reference Scheme</i> | 61.17% | 91.5% | 100% | 13.33% | 100% | 2.33% | 6.33% | 58.33% | 5.17% |
| Volume Scenario #3 | <i>Proposed Scheme</i> | 99.5% | 100% | 100% | 100% | Sleep | 36.5% | Sleep | 100% | Sleep |
| | <i>Reference Scheme</i> | 100% | 27.83% | 94.17% | 8.83% | 100% | 17.5% | 47.5% | 100% | 40.17% |

the path length between the content user and the selected content source in the PoP-level topology) are analyzed.

We first look at Fig. 3(c) to evaluate the virtual link load performance in the content distribution infrastructure. The main purpose is to check whether the additional bandwidth costs (due to the increased content traffic between PoP nodes) introduced by our scheme causes risk of congestion in the content delivery. It can be observed that our approach and the reference scheme have similar *worst-case* results in terms of maximum link load. According to the “virtual bandwidth leasing” paradigm, as long as the overall bandwidth demand does not exceed the leased bandwidth resources in the content distribution infrastructure, no traffic congestion will occur in the overlay. Therefore despite the increased virtual link loads in average, the risk of network congestion does not exist judged by the worst-case performance of link load.

Finally, we evaluate the network distances (path lengths) traversed by content requests. In general, it can be inferred that the overall path length will increase with a subset of content servers going to sleep, as some content requests need to be diverted to remote servers. However, it can be observed from Fig. 3(d) that our approach’s worst-case scenario does not increase significantly compared to the reference scheme’s results. Recall that the network distance is according to the accumulated IGP link weights along the paths taken by content requests, which reflects the end-to-end delay according to the IGP link setting. Fig. 3(d) shows that the maximum distance traversed by content requests among all scenarios is approximately 1.17 times the distance according to the reference scheme.

V. CONCLUSION

This paper has addressed the problem of optimizing content servers’ power consumptions in widely distributed content distribution infrastructures across multiple ISP domains. Based on the virtual network platform where content providers rent bandwidth resources from underlying ISP networks in order to provide end-to-end content delivery support, we proposed a novel approach that strategically puts a subset of content servers to sleep mode for power saving during off-peak hours, but without violating the load capacities of virtual links in the content delivery overlay or active content servers in the infrastructure. Such an objective is formulated into a nonlinear programming model which can be solved in an offline manner. The contribution of this work

is the consideration of optimized determination of server sleeping mode reconfiguration and intelligent mapping of user content requests to the remaining active servers by jointly taking into account server capacity and bandwidth availability.

We evaluate the proposed scheme through simulations based on interconnected GEANT-Abilene network topologies. The results showed that in various scenarios, our approach is capable of minimizing the number of active servers under load capacity constraints and reducing servers’ overall power consumptions by up to 62.2% compared to the reference scheme without power awareness. We also showed that the additional bandwidth costs our approach introduces do not have significant impact on network performances for end-to-end content delivery.

VI. ACKNOWLEDGEMENT

This work was partially funded by the EU FP7 COMET Project (248784).

REFERENCES

- [1] R. Krishnan *et al.*, “Moving beyond end-to-end information to optimize CDN performance”, in *IMC '09*, ACM, pp.190-201, Chicago, Illinois, USA.
- [2] V. Valancius, N. Laoutaris, L. Massouli, C. Diot and P. Rodriguez, “Greening the internet with nano data centers,” in *CoNEXT '09*, ACM, pp. 37-48, Rome, Italy.
- [3] L. Chiaraviglio and I. Matta, “GreenCoop: cooperative green routing with energy-efficient servers”, in *e-Energy'10*, ACM, pp.191-194, Passau, Germany.
- [4] L. Chiaraviglio and I. Matta, “An energy-aware distributed approach for content and network management,” in *Computer Communications Workshops (INFOCOM WKSHPS 2011)*, IEEE, pp.337-342, April 2011.
- [5] P. Gill, M. Arlitt, Z. Li and A. Mahanti, “YouTube traffic characterization: a view from the edge”, in *IMC '07*, ACM, pp. 15-28, San Diego, California, USA.
- [6] GEANT, “GÉANT : GÉANT”, <http://archive.geant.net/>, 2011.
- [7] Abilene, “Abilene Network”, <http://www.internet2.edu/pubs/200502-IS-AN.pdf>, 2011.
- [8] J.M. Almeida, D.L. Eager, M.K. Vernon and S.J. Wright, “Minimizing delivery cost in scalable streaming content distribution systems”, *IEEE Trans. Multimedia*, vol.6, no.2, pp. 356- 365, April 2004.
- [9] B. Nordman and K. Christensen, “Greener PCs for the enterprise,” *IT Professional*, vol.11, no.4, pp.28-37, July-Aug. 2009
- [10] T. Bektas, J.F. Cordeau, E. Erkut and G. Laporte, “Exact algorithms for the joint object placement and request routing problem in content distribution networks”, *Computers & Operations Research*, Volume 35, Issue 12, December 2008, Pages 3860-3884.
- [11] “LINGO 12.0 - Optimization Modeling Software for Linear, Nonlinear, and Integer Programming”, <http://www.lindo.com/index.php>, 2011.