

Towards Multi-Tenant Cache Management for ISP Networks

Maxim Claeys*, Daphne Tuncer†, Jeroen Famaey*, Marinos Charalambides†,
Steven Latré‡, Filip De Turck* and George Pavlou†

*Department of Information Technology, Ghent University - iMinds
Email: maxim.claeys@intec.ugent.be

†Department of Electronic & Electrical Engineering, University College London

‡Department of Mathematics and Computer Science, University of Antwerp - iMinds

Abstract—The decreasing cost of storage and the advent of virtualization technology can allow Internet Service Providers (ISPs) to deploy multi-tenant caching infrastructures and lease them to content producers and Content Delivery Networks (CDNs). Serving content requests directly from the ISP network does not only reduce the delivery time, but also allows the ISP to optimize the network resources by controlling the placement and routing of content items. In this paper, we introduce a multi-tenant cache management approach that significantly reduces the bandwidth utilization of ISPs networks by pro-actively allocating caching space, leased by content producers and/or CDNs, and intelligently routing content to the end users. Using real content request traces, we show that the optimal solution to this problem can increase the cache hit ratio by 70.64% while reducing the bandwidth usage by 57.17% on average, compared to a commonly used reactive cache management scheme. These results provide a benchmark for the development of novel multi-tenant cache management strategies.

I. INTRODUCTION

Internet Service Providers (ISPs) have been traditionally using their infrastructures as bit pipes to shift traffic between network edges. However, with the decreasing cost of storage modules and the evolution of virtualization technology, ISPs can go beyond the boundaries of simple connectivity services to enter the content delivery market, which is currently dominated by large Content Delivery Networks (CDNs) such as Akamai and Limelight. In addition to offering new revenue-generating services, ISPs can reduce CDN traffic crossing their domains, which can adversely affect network usage given that content distribution decisions are taken using limited information about carrier networks.

Previous work in the literature has proposed ISP-operated content delivery services [1], [2] and has also investigated various content management strategies based on the deployment of distributed storage within an ISP network [3]. More relevant to this work is the approach proposed by co-authors of this paper in [4], which involves operating a limited capacity CDN service within ISP networks. Lightweight content placement strategies were used to show that the proposed in-network caching functionality can enable ISPs to have better control over the utilization of their network resources.

The above approaches assume a single content provider/owner, which can be a limiting factor when considering the wide range of user interest in media content.

ISP operated content delivery services should be able to cache content from multiple sources in their network so that rich content catalogues can be offered to prospective users. In this paper, we advocate a scenario where the caching space in a distributed storage infrastructure maintained by ISPs can be leased by content producers, as well as CDN providers, that we collectively refer to as content providers. This allows them to bring their content closer to the end user. We propose an approach where caching functionality, encompassing storage resources and content placement strategies, is virtualized so that cache capacity can be flexibly allocated and the necessary management logic can be realized by commodity hardware components.

The objective of the approach we propose in this paper is to determine, given a set of content providers and the content popularity distribution in the network, which content items to store in the caching infrastructure, as well as their location, so that the capacity demand of each content provider is satisfied. The delivery path is also taken into account when computing a caching configuration with the objective to minimize the overall bandwidth consumption in the network. We formulate the problem as an Integer Linear Program (ILP) to derive the theoretical optimal placement and routing strategy. The proposed approach can be realized by a central management system that has a global view of the network resources and collects information on the popularity of contents. The results indicate a significant performance gap between the optimal solution and the Least Recently Used caching strategy.

II. RELATED WORK

Given that current content delivery services can adversely affect the utilization of ISP networks, some research efforts have been investigating new models and frameworks to support the interaction between ISPs and CDNs. These range from ISP-centric caching approaches (e.g. [1], [2]), which exclude CDNs from the delivery chain, to collaborative solutions (e.g. [5], [6]), which define new models of cooperation between ISPs and CDNs in order to improve content delivery performance.

More recently, the approach presented in [4] has proposed the operation of a limited capacity CDN service within ISP networks by deploying caches at the network edges. Such a service can allow ISPs to implement their own cache management strategy. In particular, the authors have investigated several strategies that can be used by ISPs to manage the

placement of content items in the different network caching locations according to user demand characteristics.

Various solutions to the content placement problem have been proposed in the literature, e.g. in [3], [7], [8], [9], which focused on intelligent techniques to replicate content across different network locations in order to better utilize network resources. While a distributed content placement strategy is proposed in [7] in the context of distributed replication groups, an autonomic cache management framework for information-centric networks is presented in [3]. Specific placement approaches have also been considered for hierarchical network infrastructures, especially in the context of IPTV [8], [9].

In parallel to the placement problem, previous research efforts (e.g. [10], [6]) have also focused on the server selection issue and have proposed new mechanisms to manage the redirection of user requests. In addition, optimal solution structures for the combined problem of object placement and assignment of requests to caching locations have been proposed in [11] and [12]. These, however, have not considered the problem of partitioning the available storage space between different content providers.

The decision on how to allocate capacity resources to different nodes has been considered in [13] and [14]. These focus on algorithms for the joint optimization of capacity allocation and object placement decisions under known topological and user demand information. More specifically, these aim at determining the placement of objects selected from a set of available content items and the proportion of the total storage capacity to be allocated at each potential caching location. Although a similar objective is considered in this paper, the proposed approach is significantly different. In contrast to our approach, these do not investigate how to route content requests. In addition, the approaches proposed in [13] and [14] do not consider a per-node capacity constraint, which may not be realistic in practice given that a node may not have enough capacity to accommodate all content items. It should also be noted that these focus on hierarchical caching infrastructures, whereas this restriction does not apply in this paper.

Finally, the research presented in [12] and [15] has investigated the effects on network performance of a policy that pre-partitions the storage capacity available at each caching location according to fixed ratios, which is used by different caching strategies.

III. SCENARIO

In this paper, we investigate a scenario where a large-scale ISP operates a limited capacity CDN service by deploying caching points within its network. We consider an ISP network topology consisting of edge nodes, representing access networks connecting multiple users in the same region. All content requests originate from these edge nodes. Furthermore, the network contains multiple core nodes, which interconnect all access networks. Caching capabilities are deployed at both at the edge nodes and core nodes.

A multi-tenant scenario is considered where the ISP leases the available caching space to multiple content providers. Each content provider leases a fixed amount of caching capacity to store part of its content catalogue. In this paper, we propose a

proactive content placement strategy to decide how to allocate the leased capacity and where to store which content. The proposed approach also determines from where to serve user requests, as well as the path used to deliver the content. The decisions take into account the popularity of content and its geographical distribution, which are acquired by a central manager responsible for the management decisions. In this way, the ISP aims to minimize its network resource usage, while simultaneously reducing the number of requests that have to be served from outside the ISP network.

For this purpose, the investigated problem is modelled as an ILP. Using this ILP, the optimal solution can be calculated. Since this solution comes at the cost of high management overhead, it serves as an upper limit on the performance that can be achieved with proactive content placement strategies.

IV. PROBLEM FORMULATION

To model the considered problem as an ILP, the network is represented by a directed graph $G = (V, E)$ with V and E respectively representing the set of nodes and links. The set of nodes contains both the nodes V_{ISP} , belonging to the ISP, and the external server nodes V_S , belonging to the content providers. V_{ISP} can further be divided in a set of core nodes V_C and edge nodes V_E . The links E can be divided into a set of links E_S , connected to a server node, and, ISP-managed links E_{ISP} , connecting core and edge nodes. For each node $n \in V$, we define a caching capacity $c_n \in \mathbb{N}^+$ and a set of incoming and outgoing links, denoted by $I_n \in E$ and $O_n \in E$, respectively. For every link $e \in E$, the available bandwidth capacity is denoted as $b_e \in \mathbb{N}^+$.

A set of content providers P lease caching space from the ISP. For each content provider $p \in P$, the leased amount of caching space and the set of offered content items are denoted by $d_p \in \mathbb{N}^+$ and O^p , respectively. Furthermore, each content provider $p \in P$ owns one or more servers $S_p \subset V_S$. $O = \bigcup_{p \in P} O^p$ represents the total set of offered content. Every content item $o \in O$ has an associated size $s_o \in \mathbb{N}^+$, bitrate $b_o \in \mathbb{N}^+$ and request rate $r_{o,d} \in \mathbb{N}$, for every edge node $d \in V_E$. The set of edge nodes requesting a given content $o \in O$ is denoted by $V^o \subset V_E$.

A solution to the problem consists of a content placement and routing strategy. This can be translated into binary decision variables $x_{n,o} \in \{0, 1\}$ defining if an ISP node $n \in V_{ISP}$ is used to store content $o \in O$ and $y_{e,d,o} \in \{0, 1\}$ equal to 1 if link $e \in E$ is used to deliver content $o \in O$ to edge node $d \in V_E$. Additionally, auxiliary decision variables $z_{n,o,d} \in \{0, 1\}$ are needed to define if a node $n \in V_{ISP}$ is used to store content $o \in O$ to be delivered to edge node $d \in V_E$. For these auxiliary variables, the constraints $z_{n,o,d} \leq x_{n,o}$ and $x_{n,o} \leq \sum_{d \in V_E} z_{n,o,d}$ should be applied to determine a valid solution.

Even though multiple other optimization criteria like cache hit ratio maximization are possible, the optimal solution to the problem is defined, in this paper, to be the one minimizing the bandwidth usage inside the ISP network as our aim is to reduce the ISP resource usage. This can be translated into minimizing the objective function presented in (1). A weighting factor α is used to define the relative importance of server link usage.

Higher values of α will result in minimizing the usage of uplink bandwidth.

$$\sum_{d \in V_E} \sum_{o \in O} \left(\sum_{e \in E_S} b_{or_{o,d}} y_{e,o,d} + \alpha \sum_{e \in E_{ISP}} b_{or_{o,d}} y_{e,o,d} \right) \quad (1)$$

Using the above notations, multiple constraints can be constructed to define the set of valid solutions. No nodes and links should be used to store and route content item o to edge nodes d that do not request that content. These constraints translate to (2) and (3). Furthermore, as defined by (4), an outgoing link of an edge node $d \in V_E$ can never be used to route content to itself.

$$\forall n \in V_{ISP}, \forall o \in O, \forall d \in V_E \setminus V^o : \quad z_{n,o,d} = 0 \quad (2)$$

$$\forall e \in E, \forall o \in O, \forall d \in V_E \setminus V^o : \quad y_{e,d,o} = 0 \quad (3)$$

$$\forall o \in O, \forall d \in V^o : \quad \sum_{e \in O_d} y_{e,d,o} = 0 \quad (4)$$

Since each server node $s \in S_p$ of provider p only contains content belonging to p , only this content can be routed over outgoing server links $e \in O_s$. This is defined by (5).

$$\forall p \in P, \forall s \in S_p, \forall e \in O_s, \forall d \in V_E, \forall o \in O \setminus O^p : \quad y_{e,d,o} = 0 \quad (5)$$

A valid solution should at most reserve the leased amount of caching space for each content provider $p \in P$, while meeting the storage and bandwidth capacity limitations. These constraints are modelled in (6), (7) and (8) respectively.

$$\forall p \in P : \quad \sum_{n \in V_{ISP}} \sum_{o \in O^p} s_o x_{n,o} \leq d_p \quad (6)$$

$$\forall n \in V_{ISP} : \quad \sum_{o \in O} s_o x_{n,o} \leq c_n \quad (7)$$

$$\forall e \in E : \quad \sum_{o \in O} \sum_{d \in V_E} b_{or_{o,d}} y_{e,d,o} \leq b_e \quad (8)$$

Finally, (9), (10) and (11) model the flow conservation constraints for the core and edge nodes.

$$\forall n \in V_C, \forall o \in O, \forall d \in V_E : \quad z_{n,o,d} + \sum_{e \in I_n} y_{e,d,o} = \sum_{e \in O_n} y_{e,d,o} \quad (9)$$

$$\forall o \in O, \forall d \in V^o : \quad z_{d,o,d} + \sum_{e \in I_d} y_{e,d,o} = 1 \quad (10)$$

$$\forall d \in V_A, \forall o \in O, \forall d' \in V^o \setminus \{d\} : \quad z_{d,o,d'} + \sum_{e \in I_d} y_{e,d',o} = \sum_{e \in O_d} y_{e,d',o} \quad (11)$$

Solving this ILP results in a storage pattern represented by the values of $x_{n,o}$ and a routing strategy represented by the values of $y_{e,d,o}$ that minimize the objective function in (1) while satisfying constraints (2) – (11).

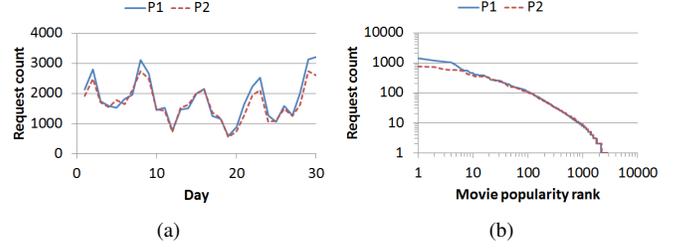


Fig. 1. Graphical representation of (a) the requests per day and (b) the popularity distribution of the evaluated VoD trace.

V. EVALUATION

A. Evaluation setup

To evaluate the proposed approach, a Video-on-Demand (VoD) scenario has been considered. For this purpose, a request trace of the VoD service of a leading European telecom operator has been used. The trace was collected over a period of 30 days between Saturday February 6, 2010 and Sunday March 7, 2010. During the considered period, 104,217 requests for 5644 unique movies were monitored, sent by 8825 unique users, originating from 12 cities. The 5644 movies were uniformly split at random between 2 content providers ($P = \{P1, P2\}$), resulting in content providers $P1$ with 2811 and $P2$ with 2833 content items. The request count per day and the popularity distribution of the content for both providers are shown in Fig. 1a and 1b, respectively. The five peaks in Fig. 1a correspond to the five weekends, with increased activity on Friday, Saturday and Sunday. In addition to this weekly pattern, clear trends can also be identified on a daily basis as well. For Wednesdays and Sundays, the activity peak is between 4:30pm and 6:30pm. For the other days of the week, most requests are reported between 8pm and 10pm. In the evaluations, the two peak hours per day are evaluated. All movies are considered to have an equal length of 90 minutes and a bitrate of 1Mbit/s ($b_o = 1 \forall o \in O$). Each movie thus has a size of 5.4Gbit ($s_o = 5400 \forall o \in O$) and is requested by the user in segments of 1 second each.

The topology used in the evaluations was based on the GÉANT topology¹, which consists of 23 nodes. Since the applied VoD request trace contains 12 cities, 12 edge nodes are needed in our topology ($V_E = \{E1, \dots, E12\}$). For both providers, one server node is assigned ($S_{P1} = \{S1\}$ and $S_{P2} = \{S2\}$). The 9 remaining nodes are modelled as core nodes ($V_C = \{C1, \dots, C9\}$). The resulting topology is shown in Fig. 2. It should be noted that the 9 most connected nodes were selected as core nodes. The links interconnecting core nodes and the links connected to the servers have a bandwidth capacity of 1Gbit/s. All other links have a capacity of 500Mbit/s. Each server hosts all the available content of the associated provider. Inside the ISP network, the storage capacity of core nodes and edge nodes is equal to 200 movies and 100 movies, respectively ($c_n = 1,080,000 \forall n \in V_C$ and $c_n = 540,000 \forall n \in V_E$). Both content providers lease storage space for 100 movies ($d_{P1} = d_{P2} = 540,000$), equal to 3.5% of their content catalogue, spread across the ISP network.

¹GÉANT Project – <http://www.geant.net>

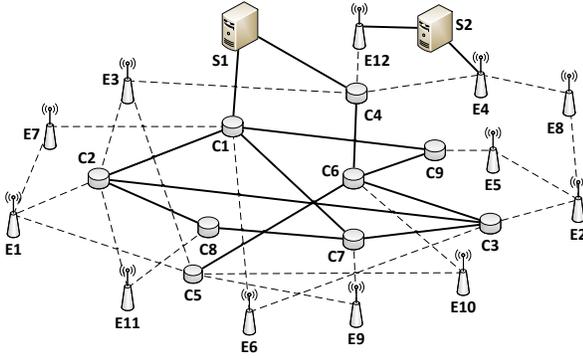


Fig. 2. Evaluated topology, based on the GÉANT topology.

B. Results

To maintain optimality throughout the evaluation period, the ILP has to be solved every time the environment changes, i.e. whenever a new request arrives or a session has finished. This optimal approach is compared to a reactive approach using the commonly used Least Recently Used (LRU) caching strategy. This approach has been performed continuously for the entire VoD trace of 1 month, while only evaluating the two peak hours of every day. For the reactive approach, the leased capacity of 100 movies for every content provider is uniformly distributed over the 21 ISP managed nodes, resulting in an available caching space equal to the size of 4.76 movies (25,714 segments of 1 second) at every node. Since content is requested at the granularity of 1 second movie segments, the reactive approach can result in only parts of a movie being available at a given node. In contrast, the proactive approach will either place an entire movie at a specific node, or not store it there at all. In the ILP, server links have been given a relative weight of $\alpha = 10$ to limit the number of requests served from outside the ISP network.

The performance of both approaches has been analyzed both in terms of cache hit ratio and bandwidth usage. Since the request pattern is different for every day of the week, results are analyzed on a per-weekday basis. As shown in Fig. 3, the average performance increase strongly fluctuates, depending on the day of the week. On average, the cache hit ratio is increased by 70.64%, from 12.97% using the LRU approach to 83.61% using the proposed proactive content placement strategy. The bandwidth usage is reduced by 57.17% on average.

To explain the fluctuating performance in terms of the day of the week, the results have been analyzed in more detail. For each of the 29 evaluated days in the VoD trace, Fig. 4 and 5 respectively show the cache hit ratio and the bandwidth consumption in terms of the number of requests for both the proactive and the reactive approach. It becomes clear from Fig. 4 that the cache hit ratio for the proactive ILP approach decreases with an increasing number of requests. This trend tends to stabilize when the number of segment requests exceeds 4 million. An increasing number of requests comes with an increased number of requested movies. Since the leased amount of caching space is fixed to 100 movies for each provider, every movie starting from the 101st has to be fetched from the origin server since it cannot be stored inside

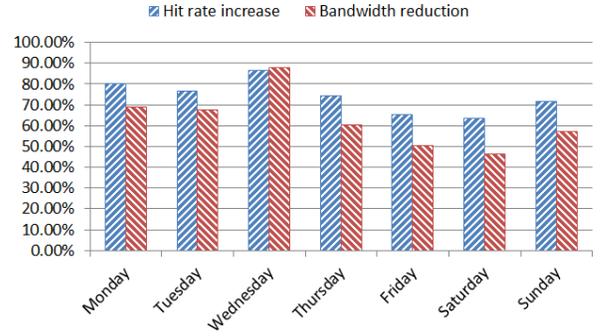


Fig. 3. Cache hit ratio increase and bandwidth reduction of the proactive ILP approach compared to the reactive LRU approach for every day of the week.

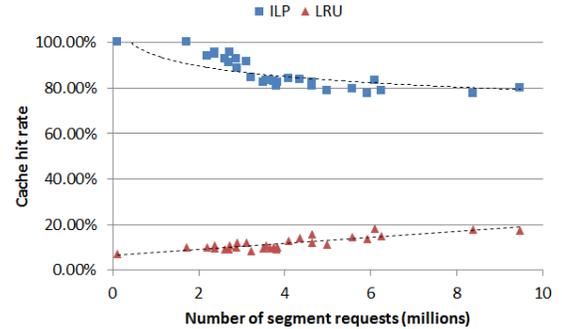


Fig. 4. Cache hit ratio for each of the evaluated days in terms of the number of requested segments.

the ISP network. At a given point, all movies are requested and the hit ratio stabilizes with an increasing number of requests. On the other hand, there is a slight increase in cache hit ratio with increasing number of requests for the reactive LRU approach. This can be explained by an increasing slope in the content popularity curve in the VoD trace when the number of requests increases. The LRU policy benefits from the fact that the most popular contents get relatively more popular when the number of requests increases. However, a large gap remains between the performance of the two approaches.

When considering the bandwidth consumption, Fig. 5 clearly shows that the proactive ILP approach outperforms the reactive LRU strategy. However, the relative performance gain decreases with an increasing number of requests. When the number of requests is low, the ILP approach is able to store the content in the cache associated with the requesting edge node, resulting in a bandwidth usage of 0. However, when the caching space is fully used, all the requests for a given movie have to be served by the origin server, resulting in longer delivery paths and higher bandwidth usage. Given that using the LRU strategy movies can be fragmented across the caching nodes, this approach is less sensitive to bandwidth increase. However, as in the case of the cache hit ratio, the two approaches have a significant performance gap.

The above observations are summarized in Fig. 6, depicting the performance gain considering cache hit ratio and bandwidth usage in terms of the number of requested segments, grouped by the day of the week.

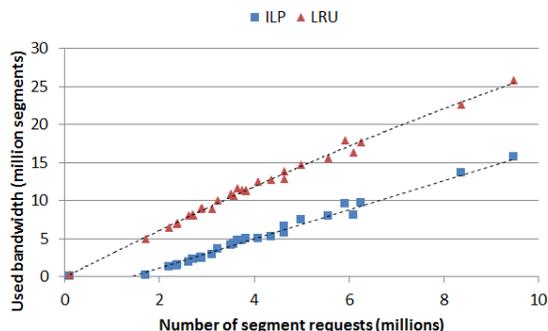


Fig. 5. Bandwidth usage for each of the evaluated days in terms of the number of requested segments.

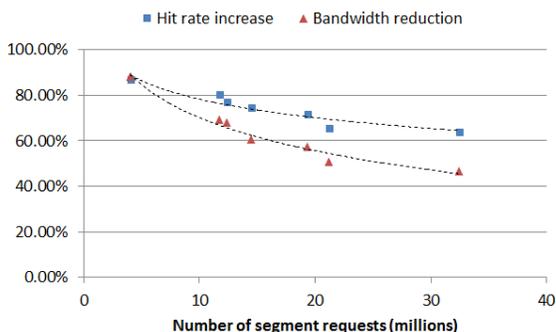


Fig. 6. Average performance increase per day of the week in terms of the number of requested segments.

C. Discussion

The results presented in Section V-B demonstrate that a reactive approach using the LRU cache replacement strategy is not well suited for a distributed caching scenario given the lack of coordinated decisions. Furthermore, we have shown that a proactive content placement strategy can significantly improve the performance, both in terms of cache hit ratio and bandwidth usage. However, since pro-actively placing content at a specific node requires the entire movie to be transferred to that node, adjusting the placement every time the request pattern changes can incur significant overhead in a practical scenario. Furthermore, this approach strongly depends on the prediction of future request patterns as optimality is severely reduced by inaccurate predictions. Nonetheless, the performance of proactive approaches is shown to be substantial. In our future work, we plan to develop efficient heuristics that also take into account the identified issues, while the results presented in this paper will be used as a benchmark.

VI. CONCLUSION

In this paper, we proposed a proactive cache management approach for ISP networks in a scenario with multiple content providers. Using a central management system that acquires knowledge about the request patterns and the popularity of content, our results suggest that the proposed approach outperforms the commonly used LRU replacement strategy. On average, the cache hit ratio can be increased from 12.97% to 83.61% while using 57.17% less bandwidth. However,

this approach comes with a high management overhead. The presented approach shows that using a reactive approach can result in situations far away from the theoretical optimum. The development of proactive approaches, considering the tradeoff between the optimality of the solution and the associated management overhead, is the subject of future research.

ACKNOWLEDGMENT

M. Claeys is funded by grant of the Agency for Innovation by Science and Technology in Flanders (IWT). This work was partly funded by FLAMINGO, a Network of Excellence project (318488) supported by the European Commission under its Seventh Framework Programme.

REFERENCES

- [1] N. Kamiyama, T. Mori, R. Kawahara, S. Harada, and H. Hasegawa, "ISP-operated CDN," in *Proc. of the IEEE INFOCOM Workshops*, 2009, pp. 49–54.
- [2] K. Cho, H. Jung, M. Lee, D. Ko, T. Kwon, and Y. Choi, "How can an ISP merge with a CDN?" *IEEE Commun. Mag.*, vol. 49, no. 10, pp. 156–162, oct. 2011.
- [3] V. Sourlas, P. Flegkas, L. Gkatzikis, and L. Tassioulas, "Autonomic cache management in information-centric networks," in *Proc. IEEE Network Operations and Management Symposium (NOMS'12)*, apr. 2012, pp. 121–129.
- [4] D. Tuncer, M. Charalambides, R. Landa, and G. Pavlou, "More control over network resources: An ISP caching perspective," in *Network and Service Management (CNSM), 2013 9th International Conference on*, Oct 2013, pp. 26–33.
- [5] W. Jiang, R. Zhang-Shen, J. Rexford, and M. Chiang, "Cooperative content distribution and traffic engineering in an ISP network," in *Proc. SIGMETRICS '09*, Seattle, WA, USA, 2009, pp. 239–250.
- [6] B. Frank, I. Poese, G. Smaragdakis, S. Uhlig, and A. Feldmann, "Content-aware traffic engineering," *SIGMETRICS Perform. Eval. Rev.*, vol. 40, no. 1, pp. 413–414, jun. 2012.
- [7] N. Laoutaris, O. Telelis, V. Zissimopoulos, and I. Stavrakakis, "Distributed selfish replication," *IEEE Trans. Parallel Distrib. Syst.*, vol. 17, no. 12, pp. 1401–1413, dec. 2006.
- [8] S. Borst, V. Gupta, and A. Walid, "Distributed caching algorithms for content distribution networks," in *Proc. IEEE INFOCOM'10*, mar. 2010, pp. 1–9.
- [9] J. Dai, Z. Hu, B. Li, J. Liu, and B. Li, "Collaborative hierarchical caching with dynamic request routing for massive content distribution," in *Proc. IEEE INFOCOM'12*, mar. 2012, pp. 2444–2452.
- [10] V. Valancius, B. Ravi, N. Feamster, and A. C. Snoeren, "Quantifying the Benefits of Joint Content and Network Routing," in *Proceedings of the ACM SIGMETRICS/International Conference on Measurement and Modeling of Computer Systems*, ser. SIGMETRICS '13, 2013, pp. 243–254.
- [11] I. Baev, R. Rajaraman, and C. Swamy, "Approximation Algorithms for Data Placement Problems," *SIAM J. Comput.*, vol. 38, no. 4, pp. 1411–1429, aug. 2008.
- [12] D. Applegate, A. Archer, V. Gopalakrishnan, S. Lee, and K. K. Ramakrishnan, "Optimal content placement for a large-scale vod system," in *Proc. Co-NEXT '10*, Philadelphia, Pennsylvania, 2010, pp. 1–12.
- [13] N. Laoutaris, V. Zissimopoulos, and I. Stavrakakis, "Joint object placement and node dimensioning for internet content distribution," *Inf. Process. Lett.*, vol. 89, no. 6, pp. 273–279, Mar. 2004.
- [14] —, "On the optimization of storage capacity allocation for content distribution," *Comput. Netw.*, vol. 47, no. 3, pp. 409–428, Feb. 2005.
- [15] A. Sharma, A. Venkataramani, and R. K. Sitaraman, "Distributing Content Simplifies ISP Traffic Engineering," in *Proceedings of the ACM SIGMETRICS/International Conference on Measurement and Modeling of Computer Systems*, ser. SIGMETRICS '13, 2013, pp. 229–242.