



Joint cache resource allocation and request routing for in-network caching services



Weibo Chu^{a,*}, Mostafa Dehghan^b, John C.S. Lui^c, Don Towsley^d, Zhi-Li Zhang^e

^aNorthwestern Polytechnical University, Xi'an, China

^bGoogle Inc., Cambridge, USA

^cThe Chinese University of Hong Kong, Hong Kong

^dUniversity of Massachusetts, Amherst, USA

^eUniversity of Minnesota, Minneapolis, USA

ARTICLE INFO

Article history:

Received 6 July 2017

Revised 23 October 2017

Accepted 28 November 2017

Available online 2 December 2017

Keywords:

cache resource allocation

cache partitioning

request routing

optimization

distributed algorithms

ABSTRACT

In-network caching is recognized as an effective solution to offload content servers and the network. A cache service provider (SP) always has incentives to better utilize its cache resources by taking into account diverse roles that content providers (CPs) play, e.g., their business models, traffic characteristics, preferences. In this paper, we study the cache resource allocation problem in a Multi-Cache Multi-CP environment. We propose a cache partitioning approach, where each cache can be partitioned into slices with each slice dedicated to a content provider. We propose a content-oblivious request routing algorithm, to be used by individual caches, that optimizes the routing strategy for each CP. We associate with each content provider a utility that is a function of its content delivery performance, and formulate an optimization problem with the objective to maximize the sum of utilities over all content providers. We establish the biconvexity of the problem, and develop decentralized (online) algorithms based on convexity of the subproblem. The proposed model is further extended to bandwidth-constrained and minimum-delay scenarios, for which we prove fundamental properties, and develop efficient algorithms. Finally, we present numerical results to show the efficacy of our mechanism and the convergence of our algorithms.

© 2017 Elsevier B.V. All rights reserved.

1. Introduction

In recent years, we have witnessed a dramatic increase in traffic over the Internet. It was reported that global IP traffic has grown 10 times from 2007 to 2015, and it will continue to increase three-fold by 2020 [1]. Among the various types of traffic generated by different applications, traffic from wireless and mobile devices accounts for a significant portion, i.e., according to [2] global mobile and wireless data traffic in 2016 amounted to 47 exabytes per month, that is 49% of the total IP traffic.

Current Internet faces significant challenges in serving this “Big Data” traffic. The host-to-host communication paradigm makes it rather inefficient to deliver content to geographically distributed users due to repeated transmissions of content, which results in unnecessary bandwidth wastage and prolonged user-perceived delays. The connection-oriented communication model also provides

little or poor support for user mobility – an important feature of future networks.

The overwhelming data traffic and limitations of the current Internet has led to a call for content-oriented networking solutions. Examples include CDNs (Content Delivery Networks) and ICNs [18] (Information-Centric Networks). Both advocate caching (either at network edge or network-wide) as part of network infrastructure, where content can be opportunistically cached so as to bring significant benefits such as bandwidth saving, short delays, server offloading. Due to its fundamental role in global content delivery, and the fact that cache storages are always scarce as compared to the amount of content transmitted over the Internet, how to efficiently utilize cache resources becomes a significant research topic. A flurry of recent studies focus on this area, such as modeling and characterizing caching dynamics [24,33], design and performance evaluation of caching mechanisms [4,14], to name a few.

In this paper, we envision that besides maximizing cache performance (measured in hit rate or miss probability) as most previous work concentrated, we also study how cache resources in network should be utilized in a way that better supports general

* Corresponding author.

E-mail addresses: wbchu@nwpu.edu.cn (W. Chu), cslui@cse.cuhk.edu.hk (J.C.S. Lui), towsley@cs.umass.edu (D. Towsley), zhzhang@cs.umn.edu (Z.-L. Zhang).

management purposes (e.g., QoS, fairness). Particularly, since content providers (CPs) have business relations with cache providers, a cache provider always has incentives to utilize its cache resources fully by taking into account diverse roles that CPs play in the market, e.g., their heterogeneous traffic characteristics, business models, QoS requirements. Baring this in mind, in this paper we study the problem of allocating cache resources among multiple content providers. We consider the problem in a “Multi-CP Multi-Cache” environment, where there are multiple cache resources distributed at different network locations serving user requests from multiple content providers. This is exactly the same setting for a variety of networking applications, such as CDNs, wireless/femtocell networks, web-cache design, and most recently, ICNs. Since there are multiple paths between each content provider and its end-users through caches, it naturally leads to a problem of jointly optimizing cache resource allocation and request routing. However, achieving system optimum by this joint optimization with the objectives of, e.g., maximizing network utility, poses a significant challenge since the problem is inherently combinatorial and NP-hard [7,17,32], and thus some optimization algorithms are needed to solve these problems efficiently (with low complexity) and practically (in a decentralized manner).

In this work, we propose a joint cache partitioning and cache-level content-oblivious request routing scheme, where we allow a cache provider to partition its caches into slices with each slice dedicated to a content provider, and each content provider routes its requests to caches it connects so to maximize its own utility. Note that there are two advantages of the proposed scheme: 1) cache partitioning restricts content contention for cache space into partitions for each CP, and hence it decouples the interactions among them and also provides a natural means for the cache manager to tune the performance for each CP; 2) besides its simplicity due to content-obliviousness (less state), cache-level request routing provides a unified request pattern seen by caches, which leads to nice properties, i.e., the hit probability of each content is solely affected by allocated cache amount, and the hit rate of each CP is linear in traffic volume directed to caches. Overall, our scheme is easy-to-implement and is suitable for cache resource management.

To abstract business relations between content providers and a cache provider, we associate with each CP a utility that is a function of its content delivery performance. We formulate an optimization problem in which the objective is to maximize the weighted sum of utilities over all content providers through proper cache partitioning and request routing. We prove that the formulated problem has a biconvexity structure, and hence can be effectively solved by existing algorithms [15]. We further prove that, with our proposed routing scheme, the optimal solution to the formulated problem has a special request routing configuration, i.e., all requests of each CP are directed to one cache it connects. This property together with the convexity of the resource-allocation subproblem makes it possible to design decentralized (online) algorithms to achieve optimum.

To illustrate that our model actually provides a general framework for cache resource allocation, we extend it to bandwidth-constrained and delay optimization scenarios, where there are bandwidth limitations between caches and content providers, and where the goal is to optimize content delivery latency. We formulate optimization problems for the two scenarios, and establish the same biconvexity property. In addition, we discover interesting phenomena, i.e., under bandwidth limitation the optimal solution is the one such that each CP directs its requests to at most one cache at the volume less than the maximum volume, and it either does not direct or directs requests at the maximum volume to the other caches. Based on these fundamental properties, efficient algorithms can be devised.

In summary, we make the following contributions:

- We propose a joint cache partitioning and cache-level content-oblivious request routing scheme in the context of multiple content providers and multiple caches, and formulate a utility-based optimization framework for cache resource management.
- We prove fundamental properties of the formulated problem, obtain its optimal routing structure, and then develop decentralized algorithms.
- Using utility-based framework, we further consider bandwidth-constrained and delay optimization scenarios. We formulate optimization problems for the two extensions, show that they also have nice properties which lead to efficient algorithms design.
- We perform numerical studies to validate the efficacy of our mechanism, and demonstrate convergence of the proposed decentralized algorithms to optimal solution.

The remainder of this paper is organized as follows. We review related work in Section 2. Section 3 describes problem setting and basic model. In Section 4 we formulate the joint cache resource allocation and request routing problem, prove its fundamental properties by analyzing its problem structure. In Section 5 we develop decentralized (online) algorithm for implementing utility-maximizing cache allocation. Section 7 presents numerical results and Section 8 discusses future research directions. We conclude the paper in Section 9.

2. Related work

The issue of cache resource allocation and management has been extensively studied in the context of CPU and memory caches (i.e., see [20,27] and the references therein). Clearly, the characteristics of the cache workload and problem settings are quite different from the networking environment, so that the techniques and design choices developed therein cannot be readily applied to our problem.

In the context of web caching, Kelly et al. [19] proposed a biased replacement policy for web caches to implement differentiated quality-of-service (QoS) by prioritizing cache space to servers. Ko et al. [21] presented a scalable QoS architecture for a shared cache storage which guarantees hitrates to multiple competing classes. Lu et al. [23] implemented an architecture for supporting differentiated caching services and adopted a control-theoretical approach to manage cache resources. Feldman and Chuang [9] proposed a QoS caching scheme that achieves service differentiation through preferential storage allocation and objects transitions across priority queues. A general cache partitioning model that integrates both QoS classes, content priority and popularity is also presented in [10].

In recent years, a significant research effort has been dedicated to the cache resource management issue in information-centric networks. Rossi et al. [29] proposed to allocate content storages heterogeneously across the network by considering graph-related centrality metrics. Psaras et al. [26] proposed probabilistic caching scheme and their studies suggested to put more cache resources at the network edge. Similarly, Fayazbakhsh et al. [8] demonstrated through simulations that most of performance benefits can be achieved by edge caching. Wang et al. [30] studied the problem of optimal cache resource allocation to network nodes by formulating it as a content placement problem.

Cache resource allocation among content providers in network for management purposes (e.g., QoS, fairness) is a new research topic. Araldo et al. in [3] adopted content-oblivious cache partitioning approach to maximize the bandwidth savings provided by the ISP cache for handling content encryption. While they focused on single-cache allocation, the problem we study here is in a Multi-Cache (and possibly with multiple service providers) environment. Hoteit et al. in [16] proposed a game-theoretic cache allocation ap-

proach to implement fairness among CPs. Exact traffic information of each CP is required to solve the problem. In our previous work [6], we proposed a cache partitioning approach to maximize aggregate network utilities over content providers, and demonstrated that cache partitioning actually provides performance gain as compared to sharing the cache with traditional LRU policy. In this work, we extend the problem setting to a broader Multi-Cache environment with more general cache management policies (under some mild conditions), and which unavoidably involves content request routing and cache selection issues. Therefore, the problem studied in [6] can be regarded as a special case.

To the best of our knowledge, this is the first work that addresses the joint cache resource allocation and request routing problem in the context of Multi-CP Multi-Cache network environment.

3. Problem description and basic model

3.1. Design space of in-network caching systems

When we consider in-network caching problems, two issues arise naturally: *caching policy* and *routing strategy*. The former refers to the rules how content objects are placed in cache storage, and the latter is how user requests are directed to caches. Different choices of the two parameters form a design space for in-network caching systems, and finding an appropriate design choice for the best system performance is always highly desired. Yet the problem is extremely challenging due to complex dynamics of the system. Take for example a simple caching system where there is only one cache serving user requests for content objects from one content provider (CP). It is well known that when the request stream is stationary, the best caching policy is to hold the top C most popular content objects in cache, where C is the cache size; when the request stream is non-stationary, dynamic caching policies such as LRU, FIFO, etc. are preferred.

Things get complicated when there are two or more content providers. Similar results can be obtained, i.e., static caching leads to the best system performance if the *aggregate* request stream is stationary. However, with multiple content providers, there is a new design choice with respect to how to utilize the cache resource: *partitioning vs non-partitioning (sharing)*. Unlike the way that cache resource as a whole is contended by different providers, now the cache can be divided into multiple slices and each slice can be dedicated to a content provider. In this case, a new question arises: should we partition the cache or should it be used as a whole piece of storage? If the answer is “we should partition it”, then another question: how much cache resource should be allocated to each CP for the optimal system performance? In our previous work [6] we have partially addressed these two questions. Surprisingly, it has been shown that for stationary request streams, sharing the cache with traditional policies such as LRU is statistically equivalent to partitioning it into a specific way, which makes cache sharing sub-optimal as compared to cache partitioning.

An even more complicated scenario is when there are multiple caches and multiple content providers, as shown in Fig. 1. In addition to caching policy, routing strategy becomes an important design parameter for such networks. The routing strategy can be *content-aware* and *content-oblivious*, depending on whether each CP needs to maintain the information of where content objects are served. More specifically, under content-aware routing each CP routes requests for a content to the cache that serves it, while under content-oblivious routing it simply forwards user requests to caches that are allocated to it. Obviously, routing strategy should match caching policy so as to maximize system performance, i.e., static caching with content-aware routing leads to the best system performance for stationary request streams. Furthermore, caches

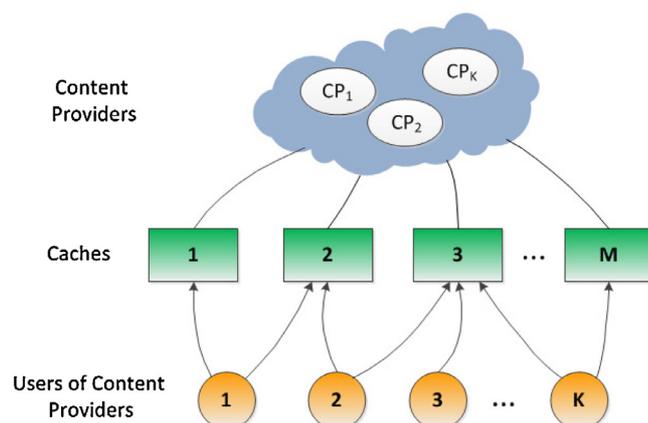


Fig. 1. A network of multiple caches and content providers.

can be *cooperative* or *non-cooperative*, where cooperative caching requires state exchange (and traffic if necessary) between caches. “Big Cache” [28], for example, abstracts the multiple caches into one single big storage and caches work collaboratively to improve system performance. Finally, it is worthy to note that other factors from real world such as delay, bandwidth constraints also play important roles in designing a practical caching system, which makes the problem even more complicated.

In this paper, instead of fully exploring the design space of in-network caching systems, we mainly focus on non-cooperative Multi-Cache Multi-CP system with cache partitioning and content-oblivious routing schemes, and address the corresponding critical design problems within a general framework.

3.2. Problem description

We now formally introduce our problem. Consider a network (e.g., a single Autonomous System) where there are M (edge) caches that serve user requests for the content from K content providers (CPs). These caches are managed by a third-party network provider, referred to as the cache manager (or cache provider) hereafter. Content providers have business relations with the cache manager and pay for cache resources. To efficiently utilize cache resources and maximize revenue, we allow the cache manager to partition its cache into multiple slices and allocate them to content providers.

Meanwhile, given cache slices, each content provider can determine the route of its user requests to these storages for its own interest. In this work, we consider cache-level content-oblivious request routing scheme, where each CP probabilistically distributes its user requests to the allocated caches. The highlights of this routing scheme is twofold: 1) it is simple and has low complexity (less state) and 2) the content popularity patterns seen by caches allocated to the same CP are identical, which further decreases the computational complexity of our proposed model. In the following section, we will show with formal proof that this routing scheme also leads to nice structural properties of the optimal solution.

Given the above setting, we further associate with each content provider a utility that is a concave, increasing function of its content delivery performance. We seek answers to the following two questions: 1) how should we partition and allocate cache resources to content providers? and 2) how should user requests be distributed to maximize the overall network cache utilization efficiency, or the weighted sum of utilities over all content providers (i.e., to implement different notion of fairness among CPs)? The problem is thus a joint cache resource allocation and request routing problem.

3.3. Cache characteristic time

Che et al. [5] introduced the notion of cache characteristic time. Based on this concept, the hit probability of a file¹ denoted as o in an LRU cache with Poisson arrivals can be approximated by

$$o(\lambda_i, T) = 1 - e^{-\lambda_i T}, \quad (1)$$

where λ_i is the request rate for file i , and T is a constant denoting the characteristic time of the LRU cache with size C . T can be computed as the unique solution to the equation

$$\sum_{i=1}^N o(\lambda_i, T) = C, \quad (2)$$

where N is the number of files in system. The cache hit rate is expressed as

$$h = \sum_{i=1}^N \lambda_i o(\lambda_i, T).$$

The characteristic time approximation has proven to be an effective tool for cache performance evaluation [13,22]. Besides LRU, it also applies to other caching policies such as FIFO, RANDOM, etc [14]. For example, FIFO and RANDOM have the same hit probability expressed as follows:

$$o(\lambda_i, T) = 1 - \frac{1}{1 + \lambda_i T}. \quad (3)$$

In this paper, we consider caching policies that can be well modeled using this characteristic time approximation. We rewrite $o(\lambda_i, T)$ as $o_i(\lambda_i, C)$ to explicitly denote that the hit probability is for file i and it is a function of the cache size C . We assume that the hit rate h is concave and increasing in C . It can be proved that this assumption holds for policies such as LRU, FIFO and Random.

Theorem 3.1. *Hit rate is a concave function of cache size for LRU, FIFO and Random policies.*

Proof. See Appendix A. \square

4. Joint cache resource allocation and request routing

In this section, we present problem formulation for the joint cache resource allocation and request routing problem. We analyze the structural property of the problem, establish its biconvexity, and give the optimal request routing strategy.

4.1. Problem formulation

Assumptions. We make the following assumptions: 1) content providers serve equal-size disjoint files; 2) requests for each file arrive according to a Poisson process; and 3) each CP k is associated with a utility $U_k(h_k)$ that is a concave, increasing function of its aggregate request hit rate h_k .

Denote by $F_k = \{f_{1k}, f_{2k}, \dots, f_{N_k k}\}$ the set of N_k different files that CP k serves, and λ_{ik} the request rate for file f_{ik} . Define $\mathbf{A} = [a_{km}]$ as a zero-one matrix denoting the connections between CPs and caches, with $a_{km} = 1$ if requests of CP k can be routed to cache m , and $a_{km} = 0$ otherwise. Request routing strategy is described by a matrix $\mathbf{P} = [p_{km}]$, with $0 \leq p_{km} \leq 1$ being the fraction of requests of CP k routed to cache m . Let C_m be the size of cache m , and C_{km} be the size of cache slice that allocated to CP k from cache m (both measured as the number of files that can be cached). Also denote

by h_{km} the request hit rate to the content of CP k that served by cache m . According to Section 3.3, we have

$$h_{km}(C_{km}, p_{km}) = \sum_{i=1}^{N_k} \lambda_{ik} p_{km} o_{ik}(\lambda_{ik} p_{km}, C_{km})$$

where $o_{ik}(\lambda_{ik} p_{km}, C_{km})$ is the hit probability, in the cache m , for file f_{ik} of CP k .

Lemma 4.1. *As long as $p_{km} \neq 0$, we have $o_{ik}(\lambda_{ik} p_{km}, C_{km}) = o_{ik}(\lambda_{ik}, C_{km})$.*

Proof. From Eqs. (1)–(3), we can see that λ_i and T appear in their product form in these equations. As a result, for the same cache size C , multiplying λ_i by p_{km} will not lead to any change in the hit probability, i.e., by letting $T = T/p_{km}$ we will have the same $1 - e^{-\lambda_i T}$ and $1 - \frac{1}{1 + \lambda_i T}$. \square

Lemma 4.1 indicates that given the cache size C , the hit probability o_{ik} is fully determined by file popularity pattern rather than the exact request volume.

Based on the above lemma, we have the following expression for h_{km} when $p_{km} \neq 0$:

$$h_{km}(C_{km}, p_{km}) = \sum_{i=1}^{N_k} \lambda_{ik} p_{km} o_{ik}(\lambda_{ik}, C_{km})$$

Observe that when $p_{km} = 0$ we have $h_{km} = 0$, and hence the above expression indeed characterizes h_{km} .

The aggregate request hit rate h_k of CP k is thus

$$h_k = \sum_{m=1}^M h_{km} = \sum_{m=1}^M \sum_{i=1}^{N_k} \lambda_{ik} p_{km} o_{ik}(\lambda_{ik}, C_{km}).$$

The joint cache resource allocation and request routing problem we study can be formulated as the following utility maximization problem:

$$\text{maximize } \sum_{k=1}^K w_k U_k(h_k(\mathbf{C}_k, \mathbf{P}_k))$$

subject to

$$\begin{aligned} \sum_{k=1}^K C_{km} &\leq C_m, \quad \forall m \\ \sum_{m=1}^M p_{km} &= 1, \quad \forall k \\ 0 &\leq p_{km} \leq a_{km}, \quad \forall k, m \end{aligned} \quad (4)$$

where $\mathbf{C}_k = (C_{k1}, C_{k2}, \dots, C_{kM})$ and $\mathbf{P}_k = (p_{k1}, p_{k2}, \dots, p_{kM})$ denote the cache allocation and request routing for CP k , respectively. $w_k > 0$ is the weight to CP k , which is chosen to reflect business preferences such as financial incentives or legal obligations. In its most simple case where $w_k = 1$ and $U_k(h_k) = h_k$, the objective becomes that of maximizing the overall cache hit rate, which provides a measure of the cache utilization efficiency.

Note that the above formulation is a “mixed-integer” programming problem that is typically hard to solve. However, in practice caches are large and therefore we assume C_{km} can take any real value, as the rounding error will be negligible.

4.2. Solution

Theorem 4.1. *The optimal solution to problem (4) is such that all requests of each CP are routed to one cache that it connects.*

Proof. See Appendix B. \square

¹ We use content and files interchangeably throughout this paper.

Theorem 4.1 states that the optimal solution to problem (4) is that requests of each CP are not splitted among caches, but rather, each CP directs all its requests to one cache it connects. This result is counter-intuitive at the first glance, since to maximize utility, each content provider will try to obtain storages from multiple caches. However, this in fact makes sense in that if this is the case, then there are multiple copies of the same files cached in these storages, which results in content duplication and hence cache resource wastage. Therefore, as a contribution of this work, here we in fact prove that cache-level content-oblivious probabilistic routing is suboptimal as compared to non-probabilistic (deterministic) request routing in a multi-cache environment.

Based on **Theorem 4.1**, Problem (4) can be reformulated as follows:

$$\begin{aligned} & \text{maximize } \sum_{k=1}^K w_k U_k(h_k(\mathbf{C}_k, \mathbf{P}_k)) \\ & \text{subject to} \\ & \sum_{k=1}^K C_{km} \leq C_m, \quad \forall m \\ & \sum_{m=1}^M p_{km} = 1, \quad \forall k \\ & p_{km} \in \{0, a_{km}\}, \quad \forall k, m \end{aligned} \quad (5)$$

Since there is a limited number of caches and CPs, problem (5) can be solved by evaluation of a series of simpler problems which are determined by the special request routing configurations given by **Theorem 4.1**. More specifically, we can convert problem (5) into a series of problems Pr_1, Pr_2, \dots, Pr_L , where L is the number of request routing configurations, and each Pr_i corresponds to a cache resource allocation problem with one specific request routing, i.e., the one that all requests of each CP are directed to one cache that it connects. The optimal solution to problem (5) is simply the one with the largest objective function value.

Theorem 4.2. *Given a request routing configuration where all requests of each CP are directed to one cache that it connects, problem (5) can be decomposed into a series of convex optimization problems, each for one cache.*

Proof. See [Appendix C](#). \square

Based on **Theorem 4.2**, we know that each Pr_i can be further divided into M subproblems, where each subproblem corresponds to a single-cache convex resource allocation problem. Let B_k be the number of caches that CP k connects. We have $L = B_1 B_2 \dots B_K$ request routing configurations. Meanwhile, solving each Pr_i involves solving M convex single-cache resource allocation problem, and hence the time complexity of the algorithm is $O(B_1 B_2 \dots B_K M)$.

The above algorithm is efficient at solving problems with small number of caches and content providers. However, as it iterates over all possible request routings, the computational time grows exponentially with the problem size. In fact, as the following theorem states, the considered problem is proved NP-hard.

Theorem 4.3. *Problem (5) is NP-hard.*

Proof. See [Appendix D](#). \square

To obtain practical solution mechanisms, we further analyze structural properties of the problem. Fortunately, the original problem turns out to be a biconvex optimization problem, as proved in **Theorem (4.4)**.

Theorem 4.4. *Problem (4) is a biconvex optimization problem.*

Proof. In **Theorem 4.2** we show that the objective function of problem (4) is concave in C_{km} 's given a fixed p_{km} 's. It is also easy to see that the objective function is concave in p_{km} 's given a fixed C_{km} 's, since h_{km} is linear in p_{km} . Therefore, the objective function is biconcave. Furthermore, as the space of the decision variables is the product of two independent convex sets, it is biconvex. As a result, problem (4) is a biconvex optimization problem. \square

The following theorem states that any solution corresponding to the special routing configurations given by **Theorem 4.1** is a *partial optimum*.

Theorem 4.5. *Let f be the objective function of problem (4). Denote by X and Y be the variable space of caching and routing, respectively. Also let y^* be a routing configuration given by **Theorem 4.1**, and $x^*(y^*)$ be the corresponding optimal caching. Then $(x^*(y^*), y^*)$ is a partial optimum, i.e., $f(x^*(y^*), y^*) \geq f(x, y^*) \forall x \in X$, and $f(x^*(y^*), y^*) \geq f(x^*(y^*), y) \forall y \in Y$.*

Biconvex optimization problems [15] have been extensively studied during the past few decades, and efficient algorithms (e.g., ACS [31] for partial optimum, GOP [12] for approximate global optimum) exist in literature. The good news from our numerical studies is that, even with local or partial optimum, the system performance can be significantly improved by the joint cache resource allocation and request routing optimization.

5. Decentralized mechanism

In this section, we present a decentralized mechanism to implement the optimal cache allocation and request routing. As compared to centralized solution, a decentralized mechanism is more desired in practice, i.e., when each CP does not want to reveal its utility function to the cache manager, or the cache manager cannot collect all traffic information of CPs. A decentralized mechanism also adapts to network changes naturally. Our proposed mechanism is based on our analysis in the above section, where the system problem can be decomposed into subproblems for the caches and for the individual content providers.

5.1. Decentralized mechanism

We have shown in **Section 4** the routing configurations of partial optimums. Furthermore, each of these routing configurations corresponds to a convex resource allocation problem. Let $CP(m)$ be the set of CPs that direct requests to cache m , and $s(k)$ be the cache to which CP k forwards all its requests. Given a fixed request routing \mathbf{P}_k 's, we have the following convex resource allocation problem:

$$\begin{aligned} & \text{maximize } \sum_{k=1}^K w_k U_k(h_k(\mathbf{C}_k)) \\ & \text{subject to} \\ & \sum_{k \in CP(m)} C_{km} \leq C_m, \quad \forall m \end{aligned} \quad (6)$$

where $U_k(h_k(\mathbf{C}_k)) = U_k(h_{ks(k)}(C_{ks(k)}, 1))$, i.e., requests of CP k are solely directed to cache $s(k)$.

The key to the decomposition of the above problem is to examine its dual. Let $\mathbf{C} = (\mathbf{C}_1, \mathbf{C}_2, \dots, \mathbf{C}_K)$ and $\boldsymbol{\lambda} = (\lambda_1, \lambda_2, \dots, \lambda_M)$. $\boldsymbol{\lambda}$ is

considered as the price vector of caches. Define Lagrangian

$$\begin{aligned} L(\mathbf{C}, \boldsymbol{\lambda}) &= \sum_{k=1}^K w_k U_k(h_k(\mathbf{C}_k)) + \sum_{m=1}^M \lambda_m (C_m - \sum_{k \in \text{CP}(m)} C_{km}) \\ &= \sum_{k=1}^K (w_k U_k(h_k(\mathbf{C}_k)) - \sum_{m=1}^M \lambda_m C_{km}) + \sum_{m=1}^M \lambda_m C_m \\ &= \sum_{k=1}^K (w_k U_k(h_k(\mathbf{C}_k)) - \lambda_{s(k)} C_{ks(k)}) + \sum_{m=1}^M \lambda_m C_m. \end{aligned}$$

The last equation holds since CP k routes all its requests to cache $s(k)$ and it will not require resources from other caches.

As the first term is separable, we have $\max_{\mathbf{C}_k} \sum_{k=1}^K$

$$(w_k U_k(h_k(\mathbf{C}_k)) - \lambda_{s(k)} C_{ks(k)}) = \sum_{k=1}^K \max_{\mathbf{C}_k} (w_k U_k(h_k(\mathbf{C}_k)) -$$

$\lambda_{s(k)} C_{ks(k)})$. According to [25], we can now readily formulate optimization problem for each CP k as follows:

CP k 's problem:

$$\max_{\mathbf{C}_k} (w_k U_k(h_k(\mathbf{C}_k)) - \lambda_{s(k)} C_{ks(k)}) \quad (7)$$

Theorem 5.1. *The cache resource allocation problem (7) has a unique optimal solution.*

Proof. Since h_{km} is concave in C_{km} , we know that h_k is concave. Now because U_k is also concave, the objective function of problem (7) is shown to be concave. Hence a unique maximizer, called the optimal solution, exists. \square

Once each cache m receives the required cache amount by content providers, its price can be adjusted as follows:

$$\lambda_m^{t+1} = [\lambda_m^t + \gamma (\sum_{k \in \text{CP}(m)} C_{km}^t - C_m)]^+ \quad (8)$$

where $\gamma > 0$ is a step size, t denotes time, and $[x]^+ = \max\{x, 0\}$.

In short, in the mechanism each content provider locally calculates its required cache amount from each cache based on their prices, and each cache adjusts its price based on the required resources from content providers. Meanwhile, as problem (7) is convex, there is no duality gap and hence the algorithm converges to the optimal solution.

It is worthy to note that the above procedure is for one specific request routing. To obtain the global optimal solution (i.e., when the problem size is small), all possible routings need to be explored. We emphasize that this process can be efficiently implemented in a parallel manner. More specifically, at each time t , each content provider locally calculates and maintains the information of (*utility, required_cache_size, routing*) and broadcasts them to the caches connected. On the other hand, with these information provided by CPs, each cache perceives all possible global routings and then after calculation it broadcasts the information of (*aggregate_utility, prices, global_routing*) back to the CPs. The optimal prices and cache allocations will be the one with the largest (weighted) aggregate utility when all procedures converge.

6. Bandwidth-constrained formulation and delay-Oriented optimization

6.1. Bandwidth-constrained formulation

Problem formulation (4) does not take into account bandwidth constraints. However, we argue that in practice, the volume of requests that each CP can forward to caches are limited, i.e., due to their bandwidth-based contracts. In this subsection, we consider such a scenario where there are bandwidth constraints between content providers and caches.

6.1.1. Problem formulation

We start by giving bandwidth constraints. Let V_{km} be the maximum volume of requests that CP k can forward to cache m . We explicitly assume that $\sum_{m=1}^M V_{km} \geq \sum_{i=1}^{N_k} \lambda_{ik}$, $\forall k$, as all requests of each CP needs to be served by the cache network. The bandwidth constraints between content providers and caches can be expressed as the following inequation:

$$\sum_{i=1}^{N_k} \lambda_{ik} p_{km} \leq V_{km}, \quad \forall k, m$$

which states that each CP cannot route requests to a cache at the volume exceeding its maximum. Incorporating this constraint into problem formulation (4), we have the bandwidth-constrained formulation as follows:

$$\maximize \sum_{k=1}^K w_k U_k(h_k(\mathbf{C}_k, \mathbf{P}_k))$$

subject to

$$\sum_{k=1}^K C_{km} \leq C_m, \quad \forall m \quad (c1)$$

$$\sum_{m=1}^M p_{km} = 1, \quad \forall k \quad (c2)$$

$$0 \leq p_{km} \leq a_{km}, \quad \forall k, m \quad (c3)$$

$$\sum_{i=1}^{N_k} \lambda_{ik} p_{km} \leq V_{km}, \quad \forall k, m \quad (c4)$$

where constraints (c1)~(c3) are exactly the same constraints as that in the original problem formulation.

For the above optimization problem, we have the following routing structure for its optimal solution.

Theorem 6.1. *The optimal solution to problem (9) is the one such that each CP directs its requests to at most one cache at the volume less than the maximum volume, and it either does not direct or directs requests at the maximum volume to the other caches.*

Proof. From Appendix B, we know that by moving requests of each CP from one cache to another which allocates a larger cache slice, the hit rate can be increased. Furthermore, since there is bandwidth limitation between each CP and each cache, when the bandwidth limitation to a cache is reached, each CP will try to route its remaining requests to another cache which allocates to it the next largest cache slice so as to maximize its hit rate, until the bandwidth limitation is reached. This process goes on and terminates when all requests of the CP are routed, with a routing configuration given by the theorem. \square

Similarly, with small problem sizes, we can devise efficient algorithm to solve problem (9) by converting it into a series of sub-problems Pr_1, Pr_2, \dots, Pr_L where each Pr_i corresponds to a cache resource allocation problem with one specific request routing, i.e., the request routing of each CP is as such that described by Theorem 6.1. Note that each Pr_i is convex as we have proved in Appendix C. The optimal solution to problem (9) is the one with the largest objective function value.

Likewise, we can prove the problem has a biconvexity structure and hence obtain its efficient solution algorithms. However, not all solutions corresponding to the routing configurations given by Theorem 6.1 are partial optimums. Instead, they can be obtained using algorithm ACS (*Alternate Convex Search*) [31] by first optimizing cache allocation variables, with starting points whose routing configurations are given by Theorem 6.1.

6.1.2. Decentralized mechanism

Based on [Theorem 6.1](#) and using the same Lagrangian-based dual decomposition technique as in [Section 5.1](#), we can formulate CP k 's problem with a given request routing \mathbf{P}_k 's as

CP k 's problem:

$$\text{maximize}_{\mathbf{C}_k} (w_k U_k(h_k(\mathbf{C}_k)) - \sum_{m \in H(k)} \lambda_m C_{km}) \quad (10)$$

where

- $U_k(h_k(\mathbf{C}_k)) = U_k(\sum_{m \in I(k)} h_{km}(C_{km}, \frac{V_{km}}{\sum_{i=1}^{N_k} \lambda_{ik}}) + h_{ks}(C_{ks}, 1 - \frac{\sum_{m \in I(k)} V_{km}}{\sum_{i=1}^{N_k} \lambda_{ik}}))$;
- $I(k)$ is defined as the set of caches to which CP k directs its requests at the maximum volume;
- s is the cache to which CP k directs its residual requests.

Note that $(I(k), s)$ is determined by \mathbf{P}_k .

Using the same reasoning as in the proof of [Theorem 5.1](#), we can prove the following theorem.

Theorem 6.2. *The cache resource allocation problem (10) has a unique optimal solution.*

The size of cache slice that CP k requires from cache m is thus:

$$\arg \{ \text{maximize}_{C_{km}, m \in I(k) \cup \{s\}} w_k U_k(h_k(\mathbf{C}_k)) - \sum_{m \in I(k) \cup \{s\}} \lambda_m C_{km} \} \quad (11)$$

Note that $C_{km} = 0$ if $m \notin I(k) \cup \{s\}$ as CP k will not require cache resources from these caches. The algorithm for each cache to update its price remains unchanged as that given by [Eq. \(8\)](#).

6.2. Latency optimization

This subsection considers user-centric problem formulation. In particular, we consider the scenario where the objective is to optimize user-perceived content delivery latency.

6.2.1. Problem formulation

Let d_{km} be the network delay of fetching content of CP k from cache m if the requested file is cached, and d_{km}^0 be the delay of fetching it from content custodians (remote servers), i.e., when the request is missed and cache m forwards it to the back-end server, downloads the file and forwards it back to the user. For each CP k , denote by t_{km} its average latency of fetching content through cache m , and t_k the overall average content delivery latency. We have

$$t_{km}(C_{km}, p_{km}) = \frac{d_{km} h_{km} + d_{km}^0 (\sum_{i=1}^{N_k} \lambda_{ik} p_{km} - h_{km})}{\sum_{i=1}^{N_k} \lambda_{ik} p_{km}}$$

and

$$t_k = \sum_{m=1}^M t_{km} p_{km} = \sum_{m=1}^M \frac{d_{km} h_{km} + d_{km}^0 (\sum_{i=1}^{N_k} \lambda_{ik} p_{km} - h_{km})}{\sum_{i=1}^{N_k} \lambda_{ik}}$$

Furthermore, associated with each CP k is a utility $U_k(t_k)$ that is a concave, decreasing function of t_k . We have the latency-optimization formulation as follows:

$$\text{maximize} \sum_{k=1}^K w_k U_k(t_k(\mathbf{C}_k, \mathbf{P}_k))$$

subject to

$$\begin{aligned} \sum_{k=1}^K C_{km} &\leq C_m, \quad \forall m \\ \sum_{m=1}^M p_{km} &= 1, \quad \forall k \\ 0 &\leq p_{km} \leq a_{km}, \quad \forall k, m \end{aligned} \quad (12)$$

Since t_{km} is linear in h_{km} , and $U_k(t_k)$ is a concave decreasing function of t_k , we can prove that problem (12) has similar properties as problem (4), i.e., [Theorem 4.1](#), [Theorem 4.2](#) and [Theorem 4.4](#) hold for problem (12). Therefore, the same centralized algorithm with the same complexity can be applied to obtain solutions.

6.2.2. Decentralized algorithm

Using Lagrangian-based dual decomposition technique as in [Section 5.1](#), we can formulate CP k 's problem with a given request routing \mathbf{P}_k 's as:

CP k 's problem:

$$\text{maximize}_{\mathbf{C}_k} (w_k U_k(t_k(\mathbf{C}_k)) - \lambda_{s(k)} C_{ks(k)}) \quad (13)$$

where we define $U_k(t_k(\mathbf{C}_k)) = U_k(t_{ks(k)}(C_{ks(k)}, 1))$, i.e., CP k routes all its requests to cache $s(k)$ and the required cache amount is such that it leads to the smallest content delivery latency.

It can be proved that problem (13) is convex and as a result, the same decentralized algorithm as in [Section 5.1](#) can be applied to obtain the optimal request routing and cache allocations.

7. Numerical studies and evaluation

In this section, we present numerical results to show: 1) the efficacy of our joint cache resource allocation and request routing mechanism and 2) the convergence of our decentralized algorithms to optimal solution.

7.1. Evaluation setup

We consider a cache network which comprises of 3 caches and 2 content providers. CP 1 connects to Cache 1 and Cache 2, while CP 2 connects to Cache 2 and Cache 3. Thus Cache 2 is shared by both content providers. The three caches adopt LRU policies and are of size $C_1 = 500$, $C_2 = 1200$, and $C_3 = 500$, respectively. The two content providers serve $N_1 = 10^4$ and $N_2 = 2 \times 10^4$ content files. File popularities for the two providers follow Zipf distributions with parameters $\alpha_1 = 0.6$ and $\alpha_2 = 0.8$, respectively. Requests for the files from the two content providers arrive as Poisson processes with aggregate rates $R_1 = 10$ and $R_2 = 15$. We set $w_1 = w_2 = 1$ assuming that the two content providers are equally important to the cache manager. The utilities of the two providers are expressed as $U_1(h) = U_2(h) = h$ so that the goal is to maximize the cache utilization efficiency (aggregate cache hit rates).

7.2. Results for basic model

7.2.1. Efficacy of the mechanism

In order to validate its efficacy, we compare the overall utility observed by our mechanism to that by simple static routing, i.e., when both content providers route their requests to the connected caches at equal probabilities. Note that even under simple static routing, cache resource allocation is obtained by solving the corresponding convex optimization problem. [Fig. 2](#) shows how the utility changes as the capacity of Cache 2 varies. It can be seen that under all cache capacities our mechanism outperforms. Moreover, the larger the shared cache, the more utility gain.

[Fig. 3](#) shows traffic distributions of the two content providers, from which we observe obvious hand-overs. More specifically, it can be seen that when the size of shared cache is small ($300 \leq C_2 \leq 500$), both content providers use its dedicated cache only; when the shared cache becomes larger ($500 < C_2 \leq 3100$), CP 2 routes all its traffic to Cache 2; and when it continues to grow ($C_2 > 3100$), both providers route all their traffic to the shared cache.

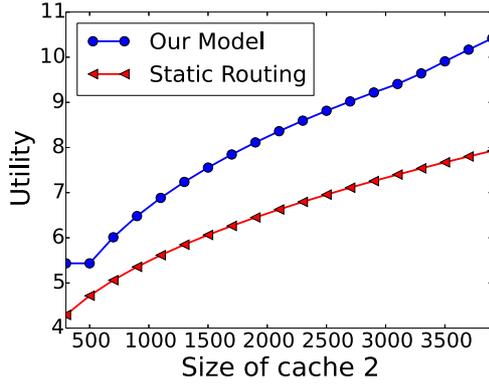


Fig. 2. Efficacy of our model as compared to static routing.

Fig. 4 shows how the shared source is allocated to the two providers. As expected, it can be seen that all the cache resource is allocated to CP 2 when only CP 2 directs requests to it. The cache resource is shared when both providers route their traffic to the cache. Note that the amount of resource allocated to them is determined by the optimization model.

From Fig. 5 we observe that the hit rate of a content provider will not keep monotonically increasing as the shared cache resource increases. In particular, we observe that the hit rate of CP 2 decreases when the capacity of Cache 2 varies from 3000 to 3100. This phenomena is due to the fact that when the size of Cache 2 is less than 3100, only CP 2 routes its traffic to Cache 2 and hence the resource is solely allocated to it; when the size of Cache 2 reaches 3100, CP 1 begins to route its traffic to it and hence the resource is shared among both providers, which certainly results in a decrease of the hit rate for CP 2. Nevertheless, the overall hit rate keeps increasing, as expected.

7.2.2. Impact of different parameters

We next look at the effect of various parameters on the competing process for the shared cache resources by the two content providers. To achieve this, we set $C_2 = 5000$ (large enough) so that both providers will have the opportunity to direct their requests to Cache 2. We fix the parameter of CP 2, and study the effect of changing the aggregate request rate R_1 of CP 1, the weight parameter w_1 , and the skewness parameter of the Zipfian file popularity distribution. Fig. 6 shows the effect of these different parameters on the observed hit rates and allocated cache amount of Cache 2 to the two providers.

(1) **Request rate:** as R_1 increases, more cache resource is allocated to CP 1, until Cache 2 is solely occupied by CP 1. Consequently, the hit rate of CP 1 increases. Note that the hit rate of CP

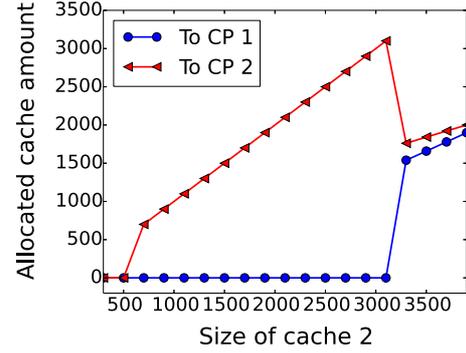


Fig. 4. Resource allocation of Cache 2 to the two content providers.

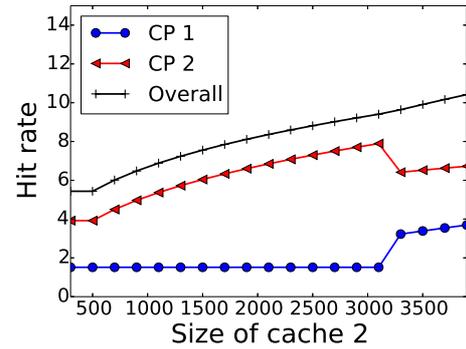


Fig. 5. Observed hit rates as a function of the shared cache size.

1 grows exactly in linear with its request rate when $R_1 > 20$ since from then on the allocated cache resource to CP 1 does not change anymore. The results indicate that if the goal is to maximize the overall hit rate, then content providers with a larger request rate will have priority over the others in the competing process.

(2) **Skewness parameter:** when α_1 increases, less cache resource is allocated to CP1. This is because with a large α_1 , a small fraction of files generates most of the traffic. Also it is surprisingly to observe that the hit rate of CP 1 increases and so does the overall hit rate. From this point of view, we conclude that a larger skewness of traffic distribution not only benefit content providers generating the traffic, but also others in system.

(3) **Weight:** Clearly, the weight of content providers significantly influences the competing process for the shared cache resources. It is observed that: i) the hit rate of CP 1 as well as cache resource allocated to it grows with the increase of its weight w_1 , until CP 1 fully occupies Cache 2; ii) although the overall hit rate also grows as we show in the figure, we argue that this is not the case in general. In fact, since there are only two content providers,

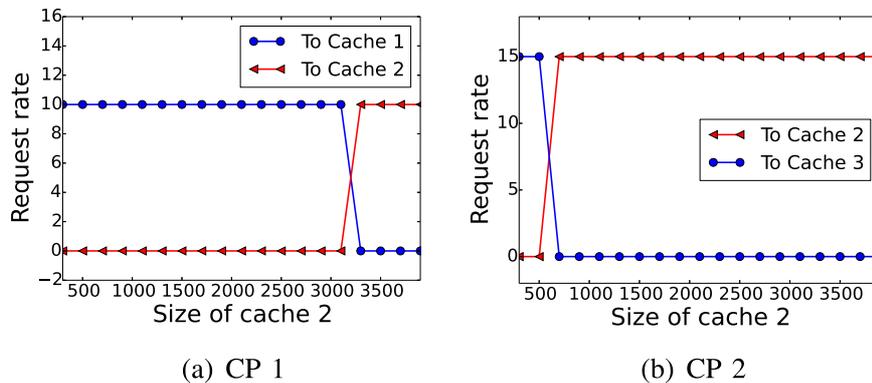


Fig. 3. Traffic distribution of the two content providers.

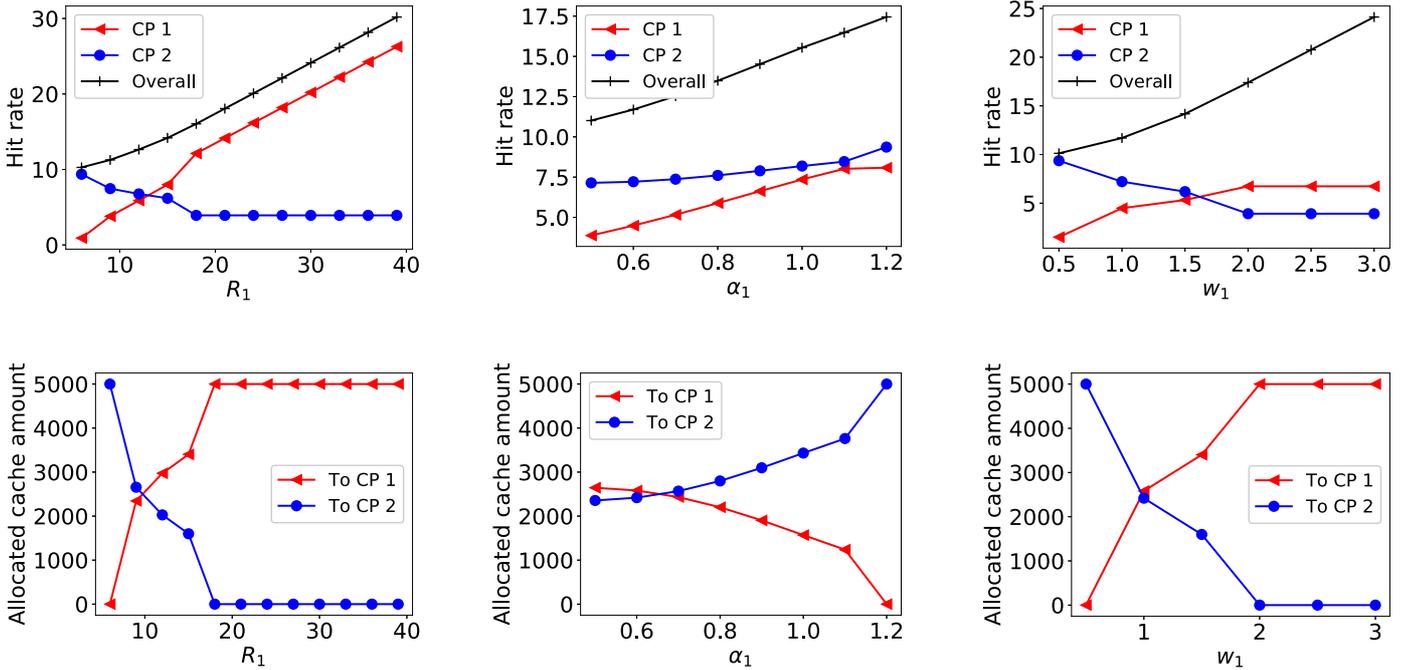


Fig. 6. Effect of the parameters on hit rates and cache resources allocated to content providers.

the figures can also be interpreted as if we change the weight of CP 2 while fixing that of CP 1. Then the hit rate of the two content providers as well as that of the system will go into the reverse direction. Indeed, as the objective function of problem (4) shows, the overall hit rate not only depends on the weights of content providers, but also depends on their traffic characteristics.

7.2.3. Convergence of decentralized algorithm

To investigate the performance of our decentralized algorithm, we fix the size of Cache 2 as $C_2 = 1900$ and choose the initial prices for the three caches as $\lambda_1^0 = \lambda_2^0 = \lambda_3^0 = 0$. The step size is set as $r = 10^{-6}$. Fig. 7 shows system dynamics as time goes on (here “CP k - Cache m ” means CP k routes all its traffic to Cache m). It can be seen that our algorithm converges under all four request routings, and the optimal solution is the one with the request routing “CP 1 - Cache 1, CP 2 - Cache2”, that is, CP 1 directs all its requests to Cache 1 and CP 2 directs all its requests to Cache 2. Looking back at Figs. 3 and 4, we can see that this result is in accordance with the centralized solution. Therefore, our algorithm converges to the optimal solution.

7.3. Bandwidth-constrained scenario evaluation

For bandwidth-constrained scenario, we set the maximum volume between CP 1 and the two caches as $V_{11} = 6$, $V_{12} = 8$, and that of CP 2 as $V_{22} = 10$ and $V_{23} = 8$. Fig. 8 shows the observed utility as the capacity of Cache 2 varies. It can be seen that even with bandwidth limitation, our mechanism outperforms static routing in this case, i.e., when both content providers route their requests to the connected caches at equal probabilities and so the traffic between them does not exceed any volume limitation.

Fig. 9 shows the optimal request routings for the two content providers, and it is interesting to observe that when Cache 2 is small ($C_2 < 500$), CP 1 routes its requests to Cache 1 at the maximum volume and CP 2 routes its requests to Cache 3 at the maximum volume. When Cache 2 becomes larger ($500 < C_2 < 1500$), CP 2 begins to route its requests to Cache 2 at the maximum volume. When it continues to grow ($C_2 > 1500$), both providers route their requests to the shared cache at the maximum volume.

Fig. 10 shows how the cache prices change under all four routings. Again we observe that our decentralized algorithm converges to the optimal solution.

We also perform numerical studies for the delay optimization scenario, and obtain similar results. To avoid redundancy and keep concise, we omit the corresponding figures. Nevertheless, we derive a consistent conclusion that our mechanism is efficient and the decentralized algorithms converge to the optimal solutions.

7.4. Performance on more complex networks

Here we report the performance of our mechanism on some more complex cache networks. In these networks, each CP connects to 2~5 caches. Content population of CPs and their corresponding Zipf parameters are set as $1 \times 10^3 \sim 5 \times 10^3$ and $0.6 \sim 0.8$, respectively. The request rate of CPs are chosen as $10 \sim 15$, and the size of caches as $200 \sim 500$. All parameters follow a uniform distribution.

For each network setting with different number of CPs and caches, we conduct numerical experiment for 20 times, and give the average observed hit rate in Table 1. The results clearly illustrate that under all three different settings, our mechanism improves system performance by 30% at least. Note that the joint optimization model is solved by SLSQP (sequential least squares programming) – a nonlinear programming solver that returns local optimum for the general NLP problem. A higher performance gain is thus expected if global algorithm, e.g., GOP [12], is applied.

Last, we emphasize that to evaluate our mechanism comprehensively it is better for us to use real-life large-scale network topologies, such as the abilene network topology (9 routers, 26 links), the dtelekom network topology (68 routers, 546 links), the geant network topology (22 routers, 66 links), to name a few. The problem of using these real-life networks is that the number of connections between content providers and caches/routers are large and grows exponentially. In that case, it makes sense to restrict content provider–cache connections within a small subset, i.e. each content provider can only route requests to caches within the neighborhood of its end-users. In that case, it suffices to de-

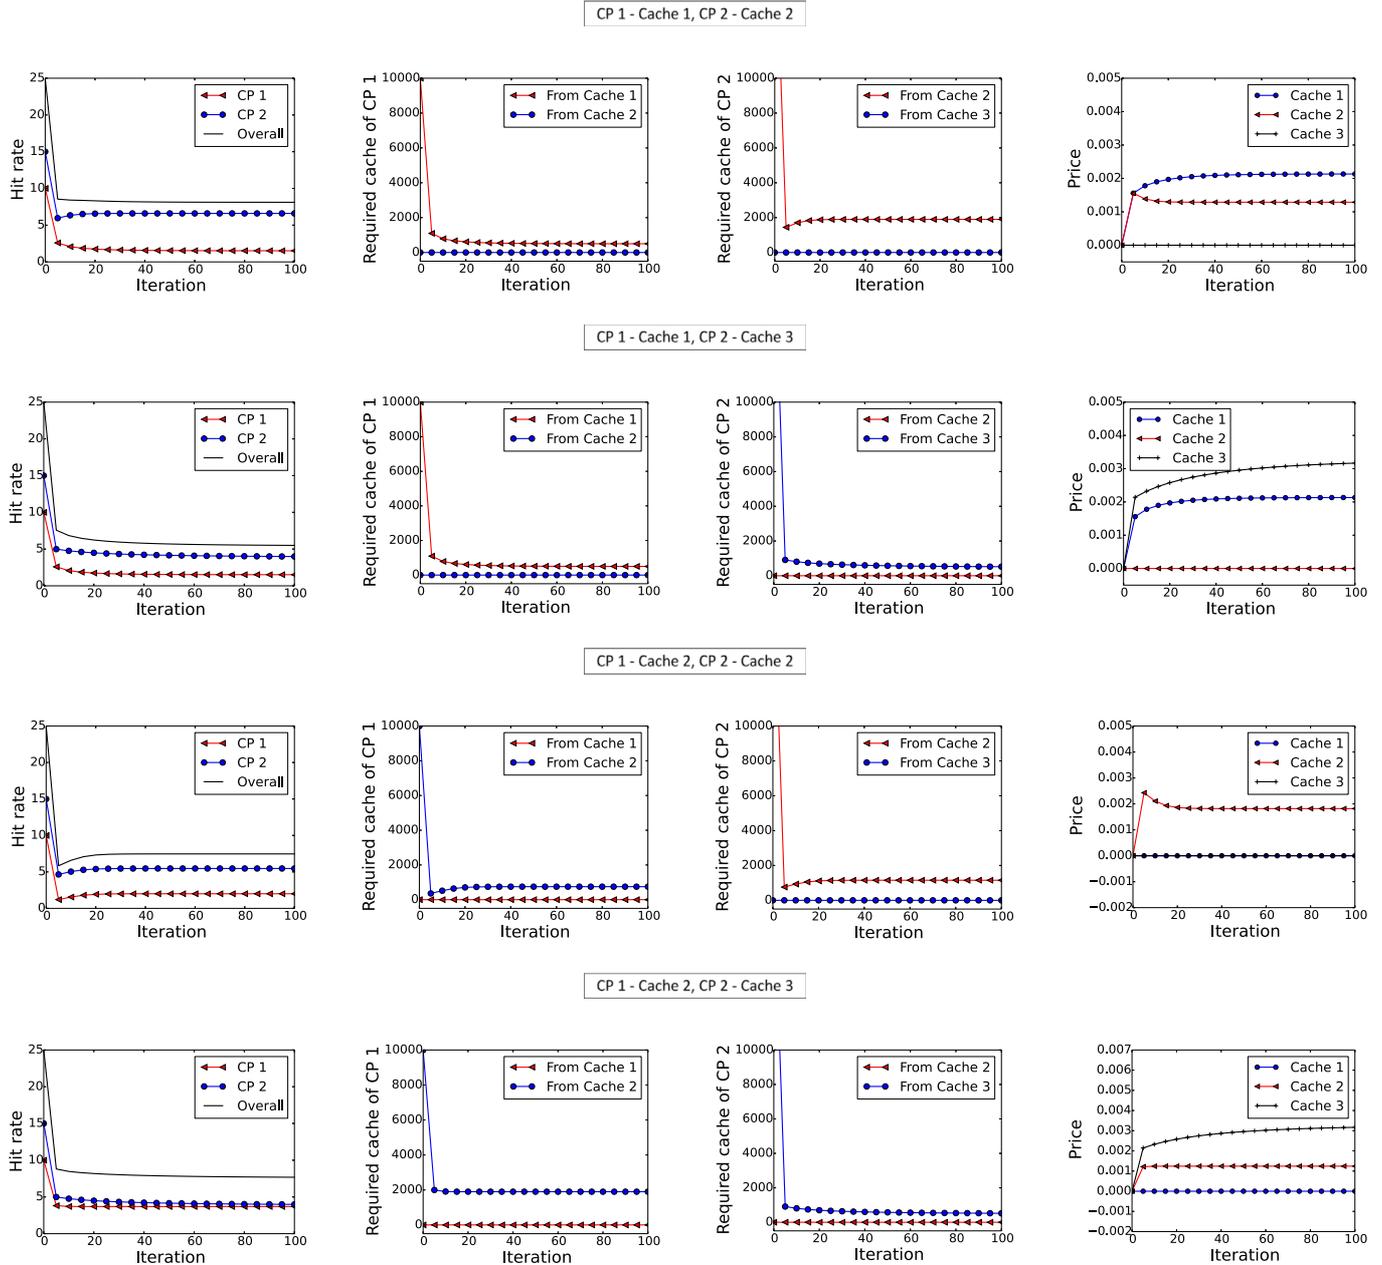


Fig. 7. Dynamics of decentralized algorithm.

Table 1
Performance on more complex networks .

# Caches	# CPs	Aggre. Hitrate by Stat. Routing	Aggre. Hitrate by Joint Opt.	Perf. Improvement
3	5	24.8	32.7	31.8%
5	10	48.6	63.4	30.4%
8	20	86.8	114.5	31.5%

compose the problem to a set of smaller sub-problems, each for a local network that is within the same order of the network that we considered in the numerical studies.

8. Discussion

In this section, we explore some implications of our framework and present some future research directions.

(1) **Non equal-size disjoint content.** Following the common practice, in this work we assume files are of equal sizes. How-

ever, files can be of variable sizes in real networks. In that case, we can divide each file into a number of small fixed-size chunks, and still treat these chunks as if they were independent disjoint files in large-scale caching systems, i.e., when the number of users accessing them are large enough. Nevertheless, this assumption needs to be carefully validated and its impact precisely measured. Another question is on the violation of the assumption of non-overlapping content, i.e., a video file can be served by multiple content providers. How to deal with these common files is another

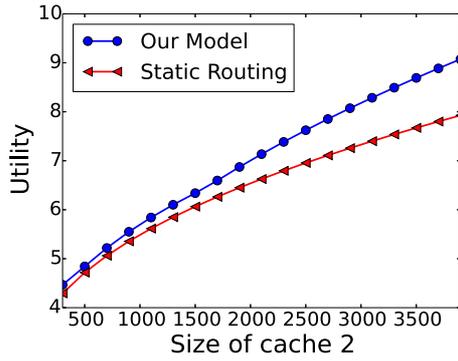
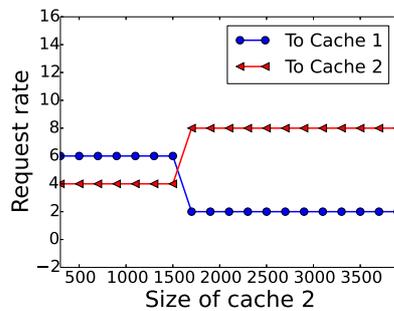


Fig. 8. Efficacy of our model with bandwidth limitation as compared to static routing.

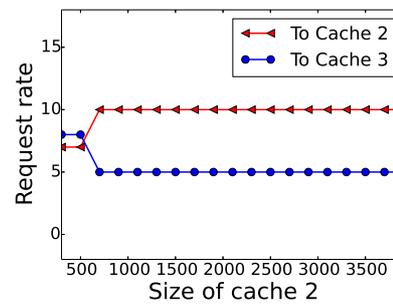
question that needs to be addressed before we apply our framework in such cases.

(2) Adaptive online algorithms. Our framework provides a means to compute optimal cache partitioning and request routing. However, it is impractical to solve the optimization problems offline and then implement the optimal policy, since in real networks, system parameters can change over time, i.e., network traffic is generally unstable. As a result, adaptive algorithms are needed. Based on the fact that network topologies/connections are more stable than traffic, here we propose a two-layer adaptive algorithm. In the upper layer, the cache manager periodically measures the connectivity of CPs and caches, the corresponding latency of content delivery, and the statistics of requests from CPs. With these measurements, the cache manager decides to recalculate request routings if there are too much change in these statistics; in the lower-layer, the cache manager collects limited traffic information for each CP and tunes the cache partitioning for them under the given routing so as to adapt to traffic changes. Of course, it remains to develop/implement such dynamic online algorithms and evaluate them in real network environment, and we leave them for our future work.

(3) Fairness. We can use different utility functions in the framework to implement different fairness among content providers. For example, the family of β -fair utility function expressed as $U_k(h_k) = \frac{h_k^{1-\beta}}{1-\beta}$ can be used to implement some interesting fairness. When $\beta \rightarrow 1$, we have $U_k(h_k) = \log h_k$, and the goal is to implement proportional fairness; when $\beta \rightarrow \infty$, it yields the objective $\max_k \min_k h_k$, which corresponds to max-min fairness. With such notions of fairness associated with utility functions, we actually provide a general and unified framework for implementing fair network resource allocation and request routing for different content providers.



(a) CP 1



(b) CP 2

Fig. 9. Traffic distribution of the two content providers with bandwidth limitation.

9. Conclusion

In this paper, we study the problem of cache resource allocation in a Multi-CP Multi-Cache environment. We propose a joint cache partitioning and content-oblivious request routing scheme, and formulate an optimization problem in which the objective is to maximize network utilities through proper cache partitioning and request routing. We give the optimal request routing strategy for each content provider, establish the biconvexity property of the formulated problem, and further develop distributed (online) algorithms. We also extend our model to the bandwidth-constrained and delay optimization scenarios to show that it provides a general framework for cache resource management in a Multi-CP Multi-Cache network. From an economic perspective, we believe that our framework also helps to build a viable market model for in-network caching services.

Acknowledgments

The work was supported by the [National Natural Science Foundation of China](#) (Grant No. 61502393), the Natural Science Basic Research Plan in Shaanxi Province of China (Grant No. 2017JM6066), the Fundamental Research Funds for the Central Universities (Grant No. 3102017zy031), and the National Science Foundation under grants CNS-1413998, CNS-1617437, CNS-1618339 and CNS-1617729. The work by John C.S. Lui is supported in part by the GRF 14200117.

Appendix A. Proof of Theorem 3.1

We rely on the characteristic time approximation for LRU, FIFO and Random policies to prove that hit rate under these policies is a concave function of the cache size.

Based on Eqs. (1) and (3), we can see that

$$\frac{\partial o(\lambda_i, T)}{\partial T} \geq 0 \quad \text{and} \quad \frac{\partial^2 o(\lambda_i, T)}{\partial T^2} \leq 0$$

For a cache of size C , we have

$$C = \sum_i o(\lambda_i, T)$$

Taking derivatives with respect to T from the two sides of the above equation, we get that

$$\frac{\partial C}{\partial T} \geq 0 \quad \text{and} \quad \frac{\partial^2 C}{\partial T^2} \leq 0$$

The above inequalities imply that C is a non-decreasing concave function of T which then means T is a non-decreasing convex function of C , i.e.,

$$\frac{\partial T}{\partial C} \geq 0 \quad \text{and} \quad \frac{\partial^2 T}{\partial C^2} \geq 0$$

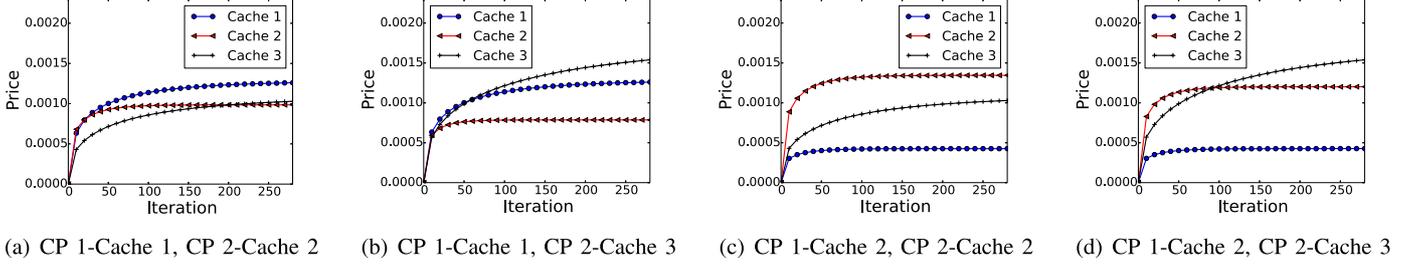


Fig. 10. Cache prices with bandwidth limitation under all four routings.

Using the equation $C = \sum_i o(\lambda_i, T)$, and taking derivatives with respect to C we obtain that

$$0 = \sum_i \frac{\partial^2 o(\lambda_i, T)}{\partial C^2}$$

We need to show

$$\sum_i \lambda_i \frac{\partial^2 o(\lambda_i, T)}{\partial C^2} \leq 0$$

Lemma A.1. Under LRU, FIFO and Random policies, for two files f_i and f_j such that $\lambda_i \geq \lambda_j$, $\frac{\partial^2 o(\lambda_j, T)}{\partial C^2} < 0$ implies that $\frac{\partial^2 o(\lambda_i, T)}{\partial C^2} < 0$.

Proof. Starting with the Eqs. (1) and (3) and taking derivatives with respect to C we obtain

$$\frac{\partial^2 o(\lambda_i, T)}{\partial C^2} = -\lambda_i^2 e^{-\lambda_i T} \left(\frac{\partial T}{\partial C} \right)^2 + \lambda_i e^{-\lambda_i T} \frac{\partial^2 T}{\partial C^2}, \quad (14)$$

for LRU policy, and

$$\frac{\partial^2 o(\lambda_i, T)}{\partial C^2} = -\frac{2\lambda_i^2}{(1 + \lambda_i T)^3} \left(\frac{\partial T}{\partial C} \right)^2 + \frac{\lambda_i}{(1 + \lambda_i T)^2} \frac{\partial^2 T}{\partial C^2}, \quad (15)$$

for FIFO and Random policies.

For LRU policy (14) implies that

$$\frac{\partial^2 o(\lambda_i, T)}{\partial C^2} < 0 \quad \text{if} \quad \frac{\partial^2 T}{\partial C^2} < \lambda_i \left(\frac{\partial T}{\partial C} \right)^2$$

On the other hand,

$$\frac{\partial^2 T}{\partial C^2} < \lambda_j \left(\frac{\partial T}{\partial C} \right)^2 \quad \text{if} \quad \frac{\partial^2 T}{\partial C^2} < \lambda_i \left(\frac{\partial T}{\partial C} \right)^2,$$

since $\lambda_j \leq \lambda_i$ and hence

$$\frac{\partial^2 o(\lambda_i, T)}{\partial C^2} < 0 \quad \text{if} \quad \frac{\partial^2 o(\lambda_j, T)}{\partial C^2} < 0.$$

Similarly, for FIFO and Random policies, (15) implies that

$$\frac{\partial^2 o(\lambda_j, T)}{\partial C^2} < 0 \quad \text{if} \quad \frac{\partial^2 T}{\partial C^2} < \frac{2\lambda_j}{1 + \lambda_j T} \left(\frac{\partial T}{\partial C} \right)^2,$$

since $\lambda_j \leq \lambda_i$ and $\frac{2\lambda_j}{1 + \lambda_j T}$ is an increasing function of λ . Therefore

$$\frac{\partial^2 o(\lambda_i, T)}{\partial C^2} < 0 \quad \text{if} \quad \frac{\partial^2 o(\lambda_j, T)}{\partial C^2} < 0. \quad \square$$

Lemma A.2. Under LRU, FIFO and Random policies, for two files f_i and f_j such that $\lambda_i \geq \lambda_j$, $\frac{\partial^2 o(\lambda_i, T)}{\partial C^2} > 0$ implies that $\frac{\partial^2 o(\lambda_j, T)}{\partial C^2} > 0$.

Lemma A.2 can be proved with similar steps as in proof of Lemma A.1.

Now, starting with the equation

$$C = \sum_i o(\lambda_i, T),$$

and taking derivative with respect to C from both sides of the equation, we obtain

$$0 = \frac{\partial^2 o(\lambda_i, T)}{\partial C^2}$$

Cache hit rate is expressed as

$$h = \sum_i \lambda_i o(\lambda_i, T)$$

and taking the derivative with respect to cache size from both sides of the equation we obtain

$$\frac{\partial h}{\partial C} = \sum_i \lambda_i \frac{\partial o(\lambda_i, T)}{\partial C} \geq 0 \quad (16)$$

Taking a second derivative yields

$$\frac{\partial^2 h}{\partial C^2} = \sum_i \lambda_i \frac{\partial^2 o(\lambda_i, T)}{\partial C^2}.$$

Since $\sum_i \frac{\partial^2 o(\lambda_i, T)}{\partial C^2} = 0$, Lemma A.1 and A.2 imply that

$$\frac{\partial^2 h}{\partial C^2} = \sum_i \lambda_i \frac{\partial^2 o(\lambda_i, T)}{\partial C^2} \leq 0 \quad (17)$$

From (16) and (17), we conclude that hit rate is an increasing concave function of the cache size.

Appendix B. Proof of Theorem 4.1

We have $h_k = \sum_{m=1}^M \sum_{i=1}^{N_k} \lambda_{ik} p_{km} o_{ik}(\lambda_{ik}, C_{km})$. Since $o_{ik}(\lambda_{ik}, C_{km})$ is increasing in C_{km} and λ_{ik} is constant, it can be seen that h_k is increased by moving requests from one cache with a small o_{ik} to the cache with a larger one, i.e., by moving requests from cache m to cache s if $o_{ik}(\lambda_{ik}, C_{km}) < o_{ik}(\lambda_{ik}, C_{ks})$. The largest value of h_k can be obtained by routing all requests of CP k to the cache with the largest o_{ik} , or equivalently, to the cache from which CP k is allocated the largest slice.

Appendix C. Proof of Theorem 4.2

Given a routing configuration, let $CP(m)$ be the set of CPs that route requests to cache m . When all requests of each CP are directed to one cache it connects, we can reformulate problem (5) as:

$$\begin{aligned} & \text{maximize} \quad \sum_{m=1}^M \sum_{k \in CP(m)} w_k U_k(h_{km}(C_{km}, 1)) \\ & \text{subject to} \quad \sum_{k \in CP(m)} C_{km} \leq C_m, \quad \forall m \end{aligned} \quad (18)$$

Since the objective function of the above problem is separable, and the variables are not coupled, problem (18) can be further decomposed into M subproblems, one for each cache. For cache m ,

the subproblem becomes:

$$\begin{aligned} & \text{maximize} \quad \sum_{k \in CP(m)} w_k U_k(h_{km}(C_{km}, 1)) \\ & \text{subject to} \quad \sum_{k \in CP(m)} C_{km} \leq C_m \end{aligned} \quad (19)$$

Lemma C.1. *The single-cache resource allocation problem (19) has a unique optimal solution.*

Proof. Since h_{km} is concave in C_{km} and U_k is also concave, the objective function in (19) is concave. Furthermore, as the feasible solution set is convex, a unique maximizer, called the optimal solution, exists. \square

Appendix D. Proof of Theorem 4.3

We have proved that the optimal solution is such that each CP directs all requests to one cache. Now if we regard bins as caches, items as CPs, the budget of bins as the capacity of caches, the weight of items as the allocated cache amount to CPs, and the profit of each item as the resulting utility of each CP, then we can see that the Separable Assignment Problem (SAP [11]) is a special case of Problem (5) (note that the weights and profits in the considered problem are not fixed, while in SAP they are fixed). SAP is a general class of maximum assignment problem with packing constraints, and it has been proved to be NP-complete since the knapsack problem is its special case [11].

References

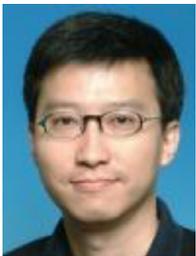
- [1] Cisco visual networking index: Forecast and methodology 2015–2020, White paper, CISCO 2016.
- [2] . The zettabyte era: Trends and analysis, CISCO, 2017. White paper
- [3] A. Araldo, G. Dan, D. Rossi, Stochastic dynamic cache partitioning for encrypted content delivery, arXiv preprint arXiv:1603.09490(2016).
- [4] G. Carofoglio, V. Gehlen, D. Perino, Experimental evaluation of memory management in content-centric networking, in: 2011 IEEE International Conference on Communications (ICC), IEEE, 2011, pp. 1–6.
- [5] H. Che, Z. Wang, Y. Tung, Analysis and design of hierarchical web caching systems, in: INFOCOM 2001. Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE, 3, IEEE, 2001, pp. 1416–1424.
- [6] W. Chu, M. Dehghan, D. Towsley, Z.-L. Zhang, On allocating cache resources to content providers, in: Proceedings of the 3rd ACM Conference on Information-Centric Networking, ACM, 2016, pp. 154–159.
- [7] M. Dehghan, B. Jiang, A. Seetharam, T. He, T. Salonidis, J. Kurose, D. Towsley, R. Sitaraman, On the complexity of optimal request routing and content caching in heterogeneous cache networks, IEEE/ACM Trans. Netw. 25 (3) (2017) 1635–1648.
- [8] S.K. Fayazbakhsh, Y. Lin, A. Tootoonchian, A. Ghodsi, T. Koponen, B. Maggs, K. Ng, V. Sekar, S. Shenker, Less pain, most of the gain: Incrementally deployable icn, in: ACM SIGCOMM Computer Communication Review, 43, ACM, 2013, pp. 147–158.
- [9] M. Feldman, J. Chuang, Service differentiation in web caching and content distribution, in: Proceedings of the IASTED International Conference on Communications and Computer Networks, 2002.
- [10] W. Feng, Y. Zhang, A general cache partition model for multiple qos classes: algorithm and simulation, in: International Conference on Computational Intelligence for Modelling, Control and Automation and International Conference on Intelligent Agents, Web Technologies and Internet Commerce (CIM-CA-IAWTIC'06), 2, IEEE, 2005, pp. 544–549.
- [11] L. Fleischer, M.X. Goemans, V.S. Mirrokni, M. Sviridenko, Tight approximation algorithms for maximum general assignment problems, in: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm, Society for Industrial and Applied Mathematics, 2006, pp. 611–620.
- [12] C.A. Floudas, V. Visweswaran, A global optimization algorithm (gop) for certain classes of nonconvex nlp-s-i, theory, Comput. Chem. Eng. 14 (12) (1990) 1397–1417.
- [13] C. Fricker, P. Robert, J. Roberts, A versatile and accurate approximation for lru cache performance, in: Proceedings of the 24th International Teletraffic Congress, International Teletraffic Congress, 2012, p. 8.
- [14] M. Garetto, E. Leonardi, V. Martina, A unified approach to the performance analysis of caching systems, ACM Trans. Model. Perform. Evaluation Comput. Syst. 1 (3) (2016) 12.
- [15] J. Gorski, F. Pfeuffer, K. Klamroth, Biconvex sets and optimization with biconvex functions: a survey and extensions, Math. Methods Oper. Res. 66 (3) (2007) 373–407.
- [16] S. Hoteit, M. El Chamie, D. Saucez, S. Secci, On fair network cache allocation to content providers, Comput. Netw. 103 (2016) 129–142.
- [17] S. Ioannidis, E. Yeh, Jointly optimal routing and caching for arbitrary network topologies, arXiv preprint arXiv:1708.05999 (2017).
- [18] V. Jacobson, D.K. Smetters, J.D. Thornton, M.F. Plass, N.H. Briggs, R.L. Braynard, Networking named content, in: Proceedings of the 5th International Conference on Emerging Networking Experiments and Technologies, ACM, 2009, pp. 1–12.
- [19] T. Kelly, Y.M. Chan, S. Jamin, J.K. MacKie-Mason, Biased replacement policies for web caches: Differential quality-of-service and aggregate user value, in: Proceedings of the 4th International Web Caching Workshop, 1999.
- [20] S. Kim, D. Chandra, Y. Solihin, Fair cache sharing and partitioning in a chip multiprocessor architecture, in: Proceedings of the 13th International Conference on Parallel Architectures and Compilation Techniques, IEEE Computer Society, 2004, pp. 111–122.
- [21] B.-J. Ko, K.-W. Lee, K. Amiri, S. Calo, Scalable service differentiation in a shared storage cache, in: Distributed Computing Systems, 2003. Proceedings. 23rd International Conference on, IEEE, 2003, pp. 184–193.
- [22] E. Leonardi, G. Luca, Least recently used caches under the shot noise model, in: 2015 IEEE Conference on Computer Communications (INFOCOM), IEEE, 2015, pp. 2281–2289.
- [23] Y. Lu, T.F. Abdelzaher, A. Saxena, Design, implementation, and evaluation of differentiated caching services, IEEE Trans. Parallel Distrib. Syst. 15 (5) (2004) 440–452.
- [24] L. Muscariello, G. Carofoglio, M. Gallo, Bandwidth and storage sharing performance in information centric networking, in: Proceedings of the ACM SIGCOMM Workshop on Information-centric Networking, ACM, 2011, pp. 26–31.
- [25] D.P. Palomar, M. Chiang, A tutorial on decomposition methods for network utility maximization, IEEE J. Sel. Areas Commun. 24 (8) (2006) 1439–1451.
- [26] I. Psaras, W.K. Chai, G. Pavlou, Probabilistic in-network caching for information-centric networks, in: Proceedings of the Second Edition of the ICN Workshop on Information-centric Networking, ACM, 2012, pp. 55–60.
- [27] M.K. Qureshi, Y.N. Patt, Utility-based cache partitioning: A low-overhead, high-performance, runtime mechanism to partition shared caches, in: Proceedings of the 39th Annual IEEE/ACM International Symposium on Microarchitecture, IEEE Computer Society, 2006, pp. 423–432.
- [28] E. Ramadan, A. Narayanan, Z.-L. Zhang, R. Li, G. Zhang, Big cache abstraction for cache networks, in: 37th IEEE International Conference on Distributed Computing Systems (ICDCS), IEEE, 2017.
- [29] D. Rossi, G. Rossini, et al., On sizing ccn content stores by exploiting topological information, in: INFOCOM Workshops, 2012, pp. 280–285.
- [30] Y. Wang, Z. Li, G. Tyson, S. Uhlig, G. Xie, Optimal cache allocation for content-centric networking, in: 2013 21st IEEE International Conference on Network Protocols (ICNP), IEEE, 2013, pp. 1–10.
- [31] R.E. Wendell, A.P. Hurter Jr, Minimization of a non-separable objective function subject to disjoint constraints, Oper. Res. 24 (4) (1976) 643–657.
- [32] E. Yeh, T. Ho, Y. Cui, M. Burd, R. Liu, D. Leong, Vip: a framework for joint dynamic forwarding and caching in named data networks, in: Proceedings of the 1st international conference on Information-centric networking, ACM, 2014, pp. 117–126.
- [33] G. Zhang, Y. Li, T. Lin, Caching in information centric networking: a survey, Comput. Netw. 57 (16) (2013) 3128–3141.



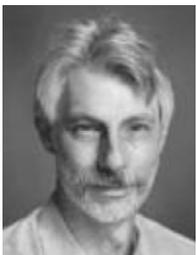
Weibo Chu received the B.S. degree in software engineering in 2005 and the Ph.D. degree in control science and engineering in 2013, both from Xi'an Jiaotong University, Xi'an, China. He has participated in various research and development projects on network testing, performance evaluation and troubleshooting, and gained extensive experiences in the development of networked systems for research and engineering purposes. From 2011–2012 he worked as a visiting researcher at Microsoft Research Asia, Beijing. From 2013 he was with the School of Computer Science and Technology, Northwestern Polytechnical University. His research interests include internet measurement and modeling, traffic analysis and performance evaluation.



Mostafa Dehghan received his Ph.D. degree from the College of Information and Computer Sciences at the University of Massachusetts Amherst. He was advised by Prof. Don Towsley, and worked on modeling and performance evaluation of content-centric networks. Cache management, routing in cache networks, and analysis of CCN data structures are some of the problems he has worked on. Before joining UMass, Mostafa received a master's degree from University of Calgary, where he developed distributed routing algorithms in cooperative networks. Mostafa is with Google now.



John C.S. Lui received the Ph.D. degree in computer science from UCLA. He is currently a professor in the Department of Computer Science and Engineering at The Chinese University of Hong Kong. His current research interests include communication networks, network/system security (e.g., cloud security, mobile security, etc.), network economics, network sciences (e.g., online social networks, information spreading, etc.), cloud computing, large-scale distributed systems and performance evaluation theory. He serves in the editorial board of *IEEE/ACM Transactions on Networking*, *IEEE Transactions on Computers*, *IEEE Transactions on Parallel and Distributed Systems*, *Journal of Performance Evaluation* and *International Journal of Network Security*. He was the chairman of the CSE Department from 2005 to 2011. He received various departmental teaching awards and the CUHK Vice-Chancellor's Exemplary Teaching Award. He is also a corecipient of the IFIP WG 7.3 Performance 2005 and IEEE/IFIP NOMS 2006 Best Student Paper Awards. He is an elected member of the IFIP WG 7.3, fellow of the ACM, fellow of the IEEE, and croucher senior research fellow.



Don Towsley holds a B.A. in Physics (1971) and a Ph.D. in Computer Science (1975) from University of Texas. From 1976 to 1985 he was a member of the faculty of the Department of Electrical and Computer Engineering at the University of Massachusetts, Amherst. He is currently a Distinguished Professor at the University of Massachusetts in the Department of Computer Science. He has held visiting positions at IBM T.J. Watson Research Center, Yorktown Heights, NY; Laboratoire MASI, Paris, France; INRIA, Sophia-Antipolis, France; AT&T Labs-Research, Florham Park, NJ; and Microsoft Research Lab, Cambridge, UK. His research interests include networks and performance evaluation. He currently serves as Editor-in-Chief of *IEEE/ACM Transactions on Networking* and on the editorial boards of *Journal of the ACM*, and *IEEE Journal on Selected Areas in Communications*, and has previously served on numerous other editorial boards. He was Program Co-chair of the joint ACM SIGMETRICS and PERFORMANCE 92 conference and the Performance 2002 conference. He is a member of ACM and ORSA, and Chair of IFIP Working Group 7.3. He has received the 2007 IEEE Koji Kobayashi Award, the 2007 ACM SIGMETRICS Achievement Award, the 1998 IEEE Communications Society William Bennett Best Paper Award, and numerous best conference/workshop paper awards. Last, he has been elected Fellow of both the ACM and IEEE.



Zhi-Li Zhang received the B.S. degree in computer science from Nanjing University, Jiangsu, China, in 1986, and the M.S. and Ph.D. degrees in computer science from the University of Massachusetts Amherst, Amherst, in 1992 and 1997, respectively. In 1997, he joined the Computer Science and Engineering faculty at the University of Minnesota, Minneapolis, MN, where he is currently a Professor. From 1987 to 1990, he conducted research with the Computer Science Department, Aarhus University, Aarhus, Denmark, under a fellowship from the Chinese National Committee for Education. He has held visiting positions with Sprint Advanced Technology Labs, Burlingame, CA; IBM T.J. Watson Research Center, Yorktown Heights, NY; Fujitsu Labs of America, Sunnyvale, CA; Microsoft Research China, Beijing, China; and INRIA, Sophia-Antipolis, France. He is a Fellow of the IEEE.