

Cache Subsidies for an Optimal Memory for Bandwidth Tradeoff in the Access Network

Original

Cache Subsidies for an Optimal Memory for Bandwidth Tradeoff in the Access Network / Ahmadi, Mahdieh; Roberts, James; Leonardi, Emilio; Movaghar, Ali. - In: IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS. - ISSN 0733-8716. - STAMPA. - 38:4(2020), pp. 736-749. [10.1109/JSAC.2020.2971806]

Availability:

This version is available at: 11583/2777985 since: 2020-04-30T08:50:06Z

Publisher:

IEEE

Published

DOI:10.1109/JSAC.2020.2971806

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2020 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Cache Subsidies for an Optimal Memory for Bandwidth Tradeoff in the Access Network

Mahdieh Ahmadi, James Roberts, Emilio Leonardi, and Ali Movaghar, *Senior Member IEEE*

Abstract—While the cost of the access network could be considerably reduced by the use of caching, this is not currently happening because content providers (CPs), who alone have the detailed demand data required for optimal content placement, have no natural incentive to use them to minimize access network operator (ANO) expenditure. We argue that ANOs should therefore provide such an incentive in the form of direct subsidies paid to the CPs in proportion to the realized savings. We apply coalition game theory to design the required subsidy framework and propose a distributed algorithm, based on Lagrangian decomposition, allowing ANOs and CPs to collectively realize the optimal memory for bandwidth tradeoff. The considered access network is a cache hierarchy with per-CP central office caches, accessed by all ANOs, at the apex, and per-ANO dedicated bandwidth and storage resources at the lower levels, including wireless base stations, that must be shared by multiple CPs.

Index Terms—Network economics, Content distribution, Caching, Cache subsidy




Fig. 1: Access network topology

I. INTRODUCTION

Since the advent of the Web and the explosive expansion of the Internet, network operators have recognized the potential for significantly reducing infrastructure costs by locally caching copies of popular contents rather than fetching them repeatedly from a remote source. However, despite the clear and increasing economic advantages of caching, it remains true that content is rarely stored at any site closer to end users than points of presence (PoPs) situated well upstream of the access network.

Network operator caching ambitions were initially stymied by more effective competition from global reach content distribution networks (CDNs) like Akamai. More significantly, content providers (CPs), like Google, Facebook and Netflix who came later and are now the source of most Internet traffic, have developed lucrative business models that rely on exclusive knowledge of their respective customer base. They preserve this knowledge by encrypting transmissions and thus prevent operators from transparently caching their contents. CPs do cache popular content in ISP PoPs both to reduce their costs and to improve end user quality of experience. However, they have hardly any additional incentive to move caches even closer to users within the access network. It is the main thesis of this paper that network operators need

to create such an incentive by financially rewarding CPs for content placements that reduce infrastructure costs. The gain from optimizing the memory for bandwidth tradeoff realized by caching is considerable and network operators, CPs and ultimately end-users will all benefit significantly.

We consider the simple tree-shaped access network topology depicted in Fig. 1. The root of the tree is a central office (CO) equipped as a datacenter [1]. This is a hub connecting multiple access network operators (ANOs) to the Internet via an ISP. An ANO might, for instance, be a mobile network operator with base stations (BSs) at the lowest layer, or a fixed network operator with optical network units (ONUs) shared by users in the same building. These ANOs may have intermediate nodes housing a cloud-radio access network (C-RAN) or an optical line terminal (OLT), respectively. These intermediate nodes are assumed to have (micro-)datacenter capabilities, like the CO, enabling flexible storage provision while base stations have just a limited capacity cache memory such as a solid-state disk. All ANOs deliver contents from all CPs and demand is assumed independent of content placements. This network model is clearly not completely general but is intended to be sufficiently generic to illustrate the challenges raised by the design and implementation of cache subsidies.

The use of subsidies to drive an optimal memory for bandwidth tradeoff in the access network is, to the authors' knowledge, an original proposal. The main contributions of the paper are as follows:

- the potential savings to be gained from an optimal trade-

M. Ahmadi and A. Movaghar are with the Department of Computer Engineering, Sharif University of Technology, Tehran, Iran. (e-mail:mahmadi@ce.sharif.edu, movaghar@sharif.edu)

J. Roberts is with the Network and Computer Science Department, Telecom ParisTech, Paris, France. (e-mail:jim.roberts@telecom-paristech.fr)

E. Leonardi is with the Politecnico di Torino, 10129 Turin, Italy (e-mail:emilio.leonardi@polito.it)

off are roughly quantified suggesting that they may count in billions of dollars per year for a large network operator;

- the need for CP subsidies as a catalyst for realizing these savings is identified as a consequence of the powerful position of the CPs in the ANO two-sided market;
- coalition game theory is applied to determine how the cost of a CP cache in the CO should be shared between the ANOs using it and how much they contribute to the subsidy;
- a distributed scheme is proposed to optimize the access network tradeoff where CPs optimally place content based on unit storage and bandwidth prices while these prices are modulated by ANOs as necessary to meet capacity constraints.

The following four sections successively present these contributions. We discuss related work in the penultimate section before drawing conclusions and highlighting directions for further study.

II. MEMORY FOR BANDWIDTH TRADEOFF

This section discusses the memory bandwidth tradeoff for an access network taking the form of a tree, assuming known demand and given storage and bandwidth costs.

A. Cache performance

We consider a catalogue of contents \mathcal{F} with a large number F of items ('files') that, to simplify but without loss of generality, we assume to be of constant size. To estimate cache hit probabilities, we assume requests for each content f generate traffic at rate λ^f content downloads per second, proportional to a popularity law q^f , i.e. $\lambda^f = Tq^f$ where $\sum_{f \in \mathcal{F}} q^f = 1$, $q^1 \geq q^2 \geq \dots \geq q^F$ and T is the overall traffic demand in downloads per second. T here is the peak traffic used for network dimensioning.

For the present work we assume the performance of a cache of size C is ideal, yielding a hit probability $h(C) = \sum_{f \leq C} q^f$. This would be realized by proactive placement of the most popular items. Note that reactive caching policies like *least recently used* (LRU) are not efficient in the access network where demand is low relative to the rate of content churn [2].

Performance depends significantly on the popularity law. Observations consistently show that request rates for the most popular items roughly follow a Zipf law, $q^f \propto f^{-\alpha}$, with $\alpha \approx 0.8$ (e.g., [3]). There is however a very large amount of content in the tail of the distribution whose popularity is imperfectly estimated but is much smaller than would be predicted by the Zipf law [4], [5]. This high volume content, like old photos posted on Facebook or family videos shared on YouTube, is stored in CP datacenters but is hardly ever cached. For the evaluations in this paper, we generally assume the catalogue \mathcal{F} does follow a Zipf law with a *useful* size F corresponding to a volume of data between 1 TB and 1 PB, it being implicitly assumed that this only represents the most popular contents.

TABLE I: Frequently Used Notation.

$k \in \mathcal{K}$	CP index
$a \in \mathcal{A}$	ANO index
$f \in \mathcal{F}$	file index
$n \in \mathcal{N}$	cache node index
$l \in \mathcal{L}$	leaf index
$i \in \mathcal{I}$	intermediate node index
\mathcal{F}_k	catalog of CP k with size $F_k = \mathcal{F}_k $
r_{ak}	CP k share of ANO a cost saving
λ_n^f	demand for content f at node n
T_a	traffic demand from ANO a users
q_a^f	popularity law of ANO a demand over \mathcal{F}
C_{nk}	cache size of CP k at node n
Λ_{nk}	residual traffic of CP k routed over link n
S_n	cache size limit at node n
B_n	bandwidth limit on link n
s_n	cost of storage at node n
b_n	cost of bandwidth between node n and its parent
$h(C)$	hit probability of cache of size C

B. Costs of storage and bandwidth

Caching trades off the cost of storage for the cost of bandwidth. While the cost of storage is well-known, it proves difficult to obtain accurate estimates of the bandwidth cost. This should be known to the network providers realizing the tradeoff but for the present paper we are obliged to make simplifying assumptions.

The cost of bandwidth is the cost of all transport and routing equipment dimensioned to carry busy period traffic between given points in the network with adequate quality of service. We assume this cost is proportional to peak demand. Costs clearly depend on the particular network instance in question but for the present discussion we consider a fixed cost per unit of busy hour demand. This might, for instance, be the long run average incremental cost (LRIC) used by telecommunications regulators.

A bandwidth cost used in prior work was derived from the marginal cost of backhaul charged to mobile network operators in France at \$2 per Mb/s per month [2]. An alternative estimate derives from the cost of bandwidth charged by Cloud providers like Google [6]. Catalogue GB per month download prices convert to a monthly rate of around \$4 per Mb/s of busy hour traffic¹.

A rough cost of storage can be deduced from the rates charged by Cloud providers. A monthly rate of around \$0.03 per GB is the current offer for frequent access storage [6].

C. Realizing the optimal tradeoff

We model an access network bringing download traffic to users as a tree network with facility for cache storage at each of its nodes (Fig. 2). The network has a set of \mathcal{N} nodes with root at node 0, the CO. Caches at the lowest tier (e.g., base stations) constitute the set of leaf nodes \mathcal{L} . Demand for content f from leaf node l is λ_l^f downloads/s. The unit cost of bandwidth between source and root is b_0 and the cost between any other node n and its parent is b_n . Unit storage cost at node n is s_n .

The placement which minimizes the total storage and bandwidth cost can be determined independently for each content.

¹Assuming busy hour demand is $1/8^{th}$ of daily traffic.




Fig. 2: Three tier access network with nodes \mathcal{N} , including root 0, leaves \mathcal{L} and intermediate nodes \mathcal{I} with storage cost s_n and bandwidth cost b_n .

To place each content f we must solve an uncapacitated facility location (UFL) problem on a tree. The UFL solution specifies the optimal set of nodes \mathcal{S} which minimizes the overall cost for the given content:

$$\mathcal{S} = \arg \min_{\mathcal{S} \subseteq \mathcal{N}} \left(\sum_{n \in \mathcal{S}} s_n + \sum_{l \in \mathcal{L}} \lambda_l^f \min_{n \in \mathcal{S}} \sum_{l \leq m < n} b_m \right), \quad (1)$$

where the summation range denoted $l \leq m < n$ covers the set of links in the unique path from l to n .

The literature provides algorithms to solve this problem, e.g. Cornuejols *et al.* [7] has an $O(N^2)$ algorithm while Shah *et al.* [8] improves this to $O(N \log N)$. The UFL algorithm should be applied successively for contents in decreasing order of overall demand, $\sum_{l \in \mathcal{L}} \lambda_l^f$. Application can cease when placement of at least one locally more popular content has previously been refused at every node.

D. Quantifying the tradeoff

While it is in general necessary to perform optimal placement using an UFL algorithm, to gain intuition we adopt a simpler network model in this subsection. We consider a symmetric 3-tier tree where demand at each leaf is statistically identical and total demand is T over all the leaves. The root at tier 3 has e_2 tier 2 children each of which has e_1 tier 1 children (the leaves).

As catalogues and cache sizes are large, it is reasonable for the present evaluation to consider F and C as real variables and to reason in terms of bytes rather than discrete contents [2]. The hit probability for a Zipf(α) popularity law is approximated by $h(C) = (C/F)^{1-\alpha}$, derived on replacing summations $\sum_{f \leq C} q^f$ by an integral. Monthly storage costs are $s^{(i)}$ at tier i nodes, i.e., a cache of size C costs $Cs^{(i)}$ per month. Bandwidth cost to bring content from the next highest tier to tier i (or from source to root when $i = 0$) is denoted $b^{(i)}$, i.e., the cost to transmit demand T is $Tb^{(i)}$ per month.

Assuming storage is possible at each node, the minimum cost network is such that tier 1 nodes cache the C_1 most popular bytes of content, tier 2 nodes cache the next C_2 most popular while the root caches the next C_3 , for some values of C_1 , C_2 and C_3 . Network cost is then,

$$\text{cost} = e_1 e_2 C_1 s^{(1)} + e_2 C_2 s^{(2)} + C_3 s^{(3)} + T \left((1 - h(C_1)) b^{(1)} + (1 - h(C_1 + C_2)) b^{(2)} + (1 - h(C_1 + C_2 + C_3)) b^{(3)} \right).$$




Fig. 3: Access network cost savings against cost factor $\Gamma = (Tb^{(i)})/(Fs^{(i)})$; $b^{(1)} = b^{(2)} = b^{(3)} = \4 per Mb/s per mth, $s^{(1)} = s^{(2)} = s^{(3)} = \0.03 per GB per mth, Zipf(0.8) popularities.

For practically relevant parameter values, standard methods yield optimal cache sizes:

$$\begin{aligned} C_1 &= F \times \min \left\{ 1, \left(\frac{(1-\alpha)Tb^{(1)}}{Fe_2(e_1s^{(1)} - s^{(2)})} \right)^{1/\alpha} \right\}, \\ C_1 + C_2 &= F \times \min \left\{ 1, \left(\frac{(1-\alpha)Tb^{(2)}}{F(e_2s^{(2)} - s^{(3)})} \right)^{1/\alpha} \right\}, \\ C_1 + C_2 + C_3 &= F \times \min \left\{ 1, \left(\frac{(1-\alpha)Tb^{(3)}}{Fs^{(3)}} \right)^{1/\alpha} \right\}. \end{aligned} \quad (2)$$

Optimal cache sizes when storage is not possible at one or two tiers can be similarly derived. Fig. 3 plots comparative gains for a network with parameter values, given in the caption, that are meant to be representative of a mobile access network [2]. Savings are plotted against the cost factor $\Gamma = (Tb^{(i)})/(Fs^{(i)})$, that here summarizes the impact of the individual parameters. The figure compares savings possible with caches at all tiers, caches at tiers 1 and 3, caches at tiers 1 and 2 and caches at tier 1 only.

A first remark is that the savings are potentially very large, especially for high Γ . For instance, a total demand volume of 10 Gb/s from 1000 BSs for a CP catalogue volume of 10 TB corresponds here to $\Gamma = 133$ and an optimal saving of more than 70% for the data of Fig. 3a. Data gathered by Perillo *et al.* [9] provide a reality check on the size of the savings possible through an optimal tradeoff: the annual revenue of network operators is measured in tens of billions of dollars (e.g., \$130B for AT&T in 2013) and CapEx is roughly 15% of this. Since most expenditure is for the access network, we estimate that potential annual savings are measured in billions of dollars.

Fig. 3a shows that most savings in this example come from caches at the two lower tiers, 1 and 2, while tier 1 BS caches alone are generally inadequate. When fanout values are inverted in Fig. 3b, however, a cache at tier 3 (the CO) brings greater savings than a cache at tier 2 for $\Gamma < 10^2$. In general, caches at all three layers significantly contribute to the overall savings.

III. THE CONTENT DELIVERY BUSINESS

We discuss the business environment arguing ANOs must subsidize CPs in order to realize the optimal memory for bandwidth tradeoff.

A. Content providers

Major content providers like Netflix, YouTube and Facebook currently generate the vast majority of Internet traffic [10]. These CPs represent multi-billion dollar businesses gaining considerably more revenue and having greater market power than the network operators on which they rely for content delivery. We include CDNs like Akamai among the CPs as they also manage large volumes of traffic and have similar business relations with the ANOs.

The CP business model generally relies on exclusive and deep knowledge of customer behavior, used for recommendation systems, ad placement and other strategic marketing activities. Customer behavior is also rightly considered to be highly confidential and for this reason alone would not be shared with any network operator seeking to realize the memory bandwidth tradeoff. The significant gains highlighted in the previous section must be realized therefore by exploiting the CPs extensive knowledge of how customer demand is distributed over its content catalogue while preserving the exclusivity of this knowledge.

B. A two-sided market

An Internet access network, viewed as a platform, is a two-sided market where end-users constitute the “money side” and CPs the “subsidy side”: end-users consume content and pay the ANO for connectivity while CPs supply the content and typically do not pay anything to the ANO for the traffic this generates [11]. The absence of compensation from the CPs is frequently the source of ANO complaints and has generated heated discussions about network neutrality (e.g., Comcast versus Netflix in the US, Free versus Google in France). To incite compensation, the ANOs might be tempted to provide less than adequate bandwidth to carry generated traffic by non-paying CPs but the impact on quality mainly hurts their paying end-user customers. CPs typically do pay the ISPs to which their servers connect but these are usually distinct from the considered ANOs. Proposals for sharing the revenue between “CP ISPs” and “eyeball ISPs” have, to our knowledge, never been implemented [12].

Of course, CPs do employ caching, typically placing dedicated servers in one or a small number of PoPs in the ISP network upstream of the CO and the access networks considered here [13]–[15]. This is advantageous both to reduce the load of their datacenters and to improve customer QoE through lower latency. The propagation time is typically much smaller from the PoP than from the remote origin datacenter and has a significant impact on latency which is in turn very important for customer satisfaction. On the other hand, the small additional reduction in propagation time due to moving the cache from the PoP to any of the access network cache locations considered here has a negligible impact on perceived latency and QoE. We believe enhanced QoE and its supposed impact on market share is therefore unlikely to motivate CP cooperation in realizing the optimal memory for bandwidth tradeoff.

C. Subsidized content placement

We claim ANOs need to persuade CPs to optimally place content in access network caches by sharing the resulting gain with them in the form of a direct subsidy. Consider, for illustration, an isolated cache dedicated to a given CP. The CP places content set C in the cache at total cost $C \times s$ where $C = |C|$. End-users generate download traffic T of which a proportion $h(C)$ is served by the cache. Bandwidth has unit cost b so that the cache brings a net saving,

$$E(C) = Th(C) \cdot b - C \cdot s. \quad (3)$$

The CP will have an incentive to cache the contents that maximize these savings if it receives a subsidy that is proportional to $E(C)$. Both ANO and CP gain with this proposal. The remainder of this paper is about how such a subsidy might be realized in practice for the hierarchical access network discussed in Sec. II-C. We first introduce the network model.

D. Network model

The considered network topology is as depicted in Fig. 1. It is a tree rooted on the central office, CO. The CO cache is shared by all the ANO instances connected to it while caches in downstream nodes, including base stations, are dedicated to their particular ANO instance. This topology is based on European Internet access. It is not perfectly general but usefully illustrates the main issues to be resolved. Extensions would be necessary, for instance, to account for ANO infrastructure sharing [16], or to more efficiently use base station caches when coverage areas overlap [17].

Multi-access edge computing (MEC) would be realized by the CO, intermediate nodes (CRANs or OLTs, say) or leaf nodes (base stations or ONUs), depending in particular on the latency requirements of the application in question. However, for *content delivery*, the latency from a cache in any location would be sufficiently small and is not a placement criterion.

We suppose the CO and transit ISP are owned by entities distinct from the ANOs, if necessary by imposed unbundling regulations. The former sell storage capacity and bandwidth at fixed unit rates and capacity is assumed unlimited. Other resources are assumed to belong to the specific ANO and their usage can be controlled by adjusting the “prices” b and s used to compute the CP subsidy.

IV. SHARING A CENTRAL OFFICE CACHE

We apply game theory to analyse value sharing between a CP and a set of ANOs who contribute to the cost of the CP cache at the central office.

A. A coalition game

In our network model in Sec. III-D, the CO hub is common to multiple ANOs. It provides cache space to CPs and is a gateway for IP transit. We analyse the memory for bandwidth tradeoff at the CO assuming ANOs are charged a common rate for transit bandwidth equivalent to b per unit of peak traffic, expressed here in content downloads per second, and the CO

storage charge is set to s per content. Available bandwidth and storage is unlimited and the values of b and s are fixed.

The cache used by each CP is distinct from that used by any other CP. This is required to preserve the privacy of customers and to ensure the integrity of the CP business models. We therefore consider a single generic CP providing content to multiple ANOs who must share the burden of the CP subsidy (cf. Sec. III-C).

To determine how savings due to caching should be shared, we apply techniques from coalition game theory. The outcome of game $G(\mathcal{N}, v)$ between players $n \in \mathcal{N}$ is determined by the value function $v(S)$ defined as the collective gain realizable by any sub-coalition $S \subseteq \mathcal{N}$. The objective is to specify how the value should be distributed between players to maximize the gain while satisfying certain properties. The value distribution, $\Phi : G \rightarrow \mathbb{R}^{|\mathcal{N}|}$, defines the amount allocated to each player. The following are two important concepts from coalition game theory [18]:

- the *core* of game $G(\mathcal{N}, v)$ is the set of all distributions Φ_n , for $n \in \mathcal{N}$, such that no group of players would gain more by leaving to form a sub-coalition $S \subset \mathcal{N}$,
- the *Shapley value* is the unique value distribution that satisfies four properties known as *efficiency*, *symmetry*, *linearity* and *null player*.

Denote the generic CP by CP and the set of ANOs by \mathcal{A} . We consider games over players $\{\text{CP} \cup \mathcal{A}\}$ with value function such that $v(S) = 0$ for any $S \subseteq \mathcal{A}$ (i.e., any sub-coalition excluding CP) and

$$v(\text{CP} \cup S) = \max_{C \subseteq \mathcal{F}} E(C, S),$$

where $E(C, S)$ is the overall saving (cf. Sec. III-C),

$$E(C, S) = \sum_{a \in S} (T_a h_a(C) \cdot b) - C \cdot s, \quad (4)$$

with $h_a(C) = \sum_{f \in C} q_a^f$.

Let $C^*(S)$ denote the set of contents that maximize savings (4) and write λ_a^f for $T_a q_a^f$, the residual demand routed to the CO for content f from ANO a users. The following proposition characterizes C^* .

Proposition IV.1.

$$C^*(S) = \left\{ f \in \mathcal{F} : \sum_{a \in S} \lambda_a^f > s/b \right\}. \quad (5)$$

Proof. We can re-write (4) as

$$E(C, S) = \sum_{f \in C} \left(\sum_{a \in S} (\lambda_a^f \cdot b) - s \right). \quad (6)$$

Clearly, $E(C, S)$ is maximized by including in C all contents for which the term in parentheses is positive. \square

B. Game between CP and ANO

To better understand the relation between CP and ANOs, we first assume the CP cache is dedicated to a single ANO, denoted ANO, and consider the game $G = (\{\text{CP}, \text{ANO}\}, v)$ with value function

$$v(\{\text{CP}, \text{ANO}\}) = E(C^*, \text{ANO}),$$

where $C^* = \arg \max_{C \subseteq \mathcal{F}} E(C, \text{ANO})$, and $v(S) = 0$ for any $S \subset \{\text{CP}, \text{ANO}\}$. It is straightforward to verify that the core of this game consists of shares

$$\begin{aligned} \Phi_{\text{CP}} &= r \times E(C^*, \text{ANO}) \\ \Phi_{\text{ANO}} &= (1 - r) \times E(C^*, \text{ANO}), \end{aligned}$$

for $0 \leq r \leq 1$.

Any member of the core with $r > 0$ has the essential property that CP will self-interestedly choose to cache the maximizing set of contents C^* . Moreover, any $r \in (0, 1)$ brings positive gain to both players and might be considered to define a suitable outcome.

In considering a similar tradeoff game, Douros *et al.* [19] propose to use the Shapley value which, in the present case can readily be shown to correspond to equal shares, $r = 0.5$. It is not obvious, however, that such a choice would be acceptable to either CP or ANO. We believe the appropriate value should reflect the relative bargaining power of CP and ANO (depending, for instance, on the size of their respective customer bases) and be decided by bilateral negotiation. Such negotiation determines a *weighted* Shapley value, as defined in [20].

We suppose the optimal tradeoff will be realized as follows. CP makes the cache placement C^* that maximizes its subsidy Φ_{CP} based on its private estimate of $\{\lambda_{\text{ANO}}^f\}$. Learning the placement size C^* , ANO pays the CO C^*s for storage and pays the transit ISP $T(1 - h(C^*))b$ for traffic. Note that the latter payment would be for realized traffic that might differ from the CP estimate. ANO additionally makes a side payment of $r(Th(C^*)b - C^*s)$ to CP. The correctness of payments is verifiable since both CP and ANO are aware of cache size C and are able to measure traffics T and $Th(C^*)$.

C. Game between multiple ANOs

Consider now multiple ANOs that can all download content from the same CP cache. The objective is to determine how the ANOs should share the cost of the cache and realize appropriate side payments to CP.

Define the game $G(\mathcal{A}, v)$ where the value function is $v(S) = E(C^*, S)$, given by (4) with C^* the maximizing set of contents for $S \subseteq \mathcal{A}$. This game does not explicitly include CP since we assume its share is determined by the share parameter r_a negotiated independently with each ANO, as discussed in Sec. IV-B. The value distribution over S , denoted $\{\Phi_a(C, S)\}$ defines the saving of ANO a for $a \in S$, to be shared with CP so that the overall subsidy is

$$\text{sub}(C, S) = \sum_{a \in S} r_a \Phi_a(C, S). \quad (7)$$

We claim the distribution $\{\Phi_a(C, \mathcal{A})\}$ should ideally satisfy the following properties:

- *coalition incentive*: the distribution should encourage all ANOs to share the same CP cache; this will happen when the distribution $\{\Phi_a(C, \mathcal{A})\}$ is in the core of game G ,
- *efficiency*: total distributed value should be equal to the value function, i.e., $\sum_{a \in \mathcal{A}} \Phi_a(C, \mathcal{A}) = E(C, \mathcal{A})$,

- *optimality*: the CP subsidy (7) must be proportional to $E(C, S)$ given by (4) to induce it to realize the overall optimal placement C^* ,
- *neutrality*: the CP placement C^* should be independent of the respective bargaining power of ANOs as manifested by the r_a ,
- *verifiability*: the savings distribution should be readily computable and verifiable by the ANOs and CP.

Supposing ANOs should only pay for their own traffic while sharing the cost of storage, we consider distributions of the form,

$$\Phi_a(C, S) = T_a h_a(C) \cdot b - \zeta_a(C) C \cdot s, \quad (8)$$

$$= \sum_{f \in C} \left(\lambda_a^f \cdot b - \eta_a(f) \cdot s \right), \quad (9)$$

where $\zeta_a(C)$ and $\eta_a(f)$ are fractional shares, to be defined, and such that $\sum_{f \in C} \eta_a(f) = \zeta_a(C) C$ and $\sum_{a \in S} \zeta_a(C) = 1$. The latter condition is required for efficiency.

Theorem IV.1 defines the unique distribution of this form that is valid for general traffic and arbitrary subsidy fractions.

Theorem IV.1. *Under general traffic $\{\lambda_a^f\}$ and arbitrary subsidy fractions $\{r_a\}$, the distribution $\{\Phi_a(C, S)\}$ in (9) is optimal and neutral over coalition $S \subseteq \mathcal{A}$, if and only if,*

$$\eta_a(f) = \frac{\lambda_a^f}{\sum_{n \in S} \lambda_n^f}, \forall a \in S. \quad (10)$$

Moreover, the distribution $\Phi_a(C, \mathcal{A})$ is in the core of $G(\mathcal{A}, v)$.

Proof. From (7) and (9), the CP subsidy is

$$\text{sub}(C, S) = \sum_{f \in C} \left(\sum_{n \in S} r_n \lambda_n^f \cdot b - \sum_{n \in S} r_n \eta_n(f) \cdot s \right). \quad (11)$$

This is maximal if and only if CP chooses to cache contents f such that

$$\frac{\sum_{n \in S} r_n \lambda_n^f}{\sum_{n \in S} r_n \eta_n(f)} > \frac{s}{b}. \quad (12)$$

For optimality, this condition must coincide with (5) so that,

$$\sum_{n \in S} \lambda_n^f = \frac{\sum_{n \in S} r_n \lambda_n^f}{\sum_{n \in S} r_n \eta_n(f)}, \forall f \in \mathcal{F}.$$

For neutrality, this equation must be an identity with respect to the $\{r_a\}$ yielding expression (10) on setting particular values $r_a = 1$ and $r_n = 0$ for $n \neq a$.

To prove sufficiency, using (10) in (12) gives the condition of Proposition IV.1. The shares are thus optimal. Sharing is also neutral since the set defined by Proposition IV.1 is independent of $\{r_a\}$.

To prove $\{\Phi_a(C, \mathcal{A})\}$ is in the core, we deduce from optimality that

$$\begin{aligned} v(S) = E(C^*, S) &= \sum_{f \in C^*(S)} \left(\sum_{a \in S} (\lambda_a^f \cdot b) - s \right) \\ &= \sum_{f \in C^*(S)} \sum_{a \in S} \left(\lambda_a^f \cdot b - \frac{\lambda_a^f}{\sum_{n \in S} \lambda_n^f} \cdot s \right) \\ &\leq \sum_{f \in C^*(\mathcal{A})} \sum_{a \in S} \left(\lambda_a^f \cdot b - \frac{\lambda_a^f}{\sum_{n \in \mathcal{A}} \lambda_n^f} \cdot s \right) \\ &= \sum_{a \in S} \Phi_a(C^*(\mathcal{A}), \mathcal{A}), \end{aligned}$$

where the last inequality follows from the fact that any content f included in the optimal set $C^*(S)$ is necessarily also included in $C^*(\mathcal{A})$. Thus no sub-coalition has value greater than that of the grand coalition. \square

The distribution defined in Theorem IV.1 is unfortunately not verifiable by the ANOs since $\zeta_a(C)$ depends on the demand distributions λ_a^f that are unknown to them. Proposition IV.2 defines a verifiable distribution that is optimal in some special cases.

Proposition IV.2. *Over coalition $S \subseteq \mathcal{A}$, the distribution $\{\Phi_a(C, S)\}$ defined in (8) is optimal and neutral if*

$$\zeta_a(C) = \frac{T_a h_a(C)}{\sum_{n \in S} T_n h_n(C)}, \quad (13)$$

and all ANOs have the same popularity distribution, $q_a^f = q^f$ for $a \in S$, and optimal if (13) holds and all ANOs apply the same subsidy fraction, $r_a = r$ for $a \in S$.

Proof. With Φ_a given by (8) and (13), the CP subsidy may be written,

$$\begin{aligned} \text{sub}(C, S) &= \frac{\sum_{n \in S} r_n T_n h_n(C)}{\sum_{n \in S} T_n h_n(C)} \times \\ &\quad \left(\sum_{a \in S} T_a h_a(C) \cdot b - C \cdot s \right). \quad (14) \end{aligned}$$

The term in parenthesis is the overall profit so that CP will place the optimal set C^* if the pre-factor is independent of C . This occurs if the popularity distributions are the same for all when $h_n(C) = h(C)$. In this case the distribution is optimal and neutral. The pre-factor is also constant if all the r_a are equal proving optimality. \square

The value distributions determined from Theorem IV.1 and Proposition IV.2 are different except in the special case where all popularity laws are the same. When the laws are different and fractions r_{ak} are not the same, the value distribution determined from (13) is not optimal. Moreover, the CP is required to maximize the subsidy defined by (14) which is non-trivial, especially in the context of the access network as considered in the next section. The optimal distribution determined from (10) is therefore preferable despite the impossibility for ANOs to independently verify their allocation.

We envisaged applying a distribution based on the Shapley value. One possibility is to derive the Shapley value for a




Fig. 4: Percentage difference when verifying subsidies computed using Theorem IV.1 by applying formula (13); 2 ANOs, $T_1 = 160$ Mb/s, $T_2 = 80$ Mb/s, $F = 10^7$, uncorrelated Zipf(0.8) popularities, $r_2 = 0.5$.

fixed cache size C , compute the CP subsidy based on this, and then suppose CP places the subsidy maximizing content C . This distribution does not have the optimality property, however: the CP does not have the correct incentive to optimize the overall memory for bandwidth tradeoff. To compute the Shapley value on supposing content placement is optimized for each ordered sub-coalition, on the other hand, rapidly becomes computationally intractable (as in [12], for instance).

D. Approximate verification

The value distribution determined from Theorem IV.1 has all the desirable properties except verifiability. In this section, we numerically explore the possibility of using ζ_a from (13) in (8) to perform an approximate verification. In other words, CP computes the optimal set C and distribution $\{\Phi_a(C, \mathcal{A})\}$ by applying Theorem IV.1 while each ANO estimates the realized shares using (13) and (8). There is of course, a difference between predicted demand (the λ_a^f) and realized demand (measured $T_a h_a(C)$) but we ignore this discrepancy here and consider the error arising when realized demand actually coincides with the forecast.

We consider a CP with a catalogue of $F = 10^7$ megabyte files. The fixed cost of storage at the CO is $s = \$0.03$ per GB ($\$3 \times 10^{-5}$ per content) and the cost of bandwidth is $b = \$4$ per Mb/s. Two ANOs share the CP cache and have demand $T_1 = 160$ Mb/s and $T_2 = 80$ Mb/s. The popularity law for both is Zipf(0.8) but the ranking of contents for ANO 2 is a random permutation of the ranking for ANO 1. Numerical results are the average of 10 different random permutations.

Fig. 4 plots the percentage difference between the required subsidy using (13), $\text{sub}_a(\zeta)$, and the subsidy due using (10), $\text{sub}_a(\eta)$, as a function of r_1 when $r_2 = 0.5$ is fixed. More precisely, we plot

$$\text{err}_a = (\text{sub}_a(\zeta) - \text{sub}_a(\eta)) / \text{sub}_a(\eta) \times 100$$

against r_1 , for $a = 1$ and $a = 2$, and the equivalent relative error, err_{tot} , computed for the total subsidies (14) and (11). The 95% confidence interval estimated from the 10 different ranking permutations is smaller than the line thickness.

As predicted by Prop. IV.2, the total subsidy is accurately estimated when $r_1 = r_2$. The errors for individual subsidies, however, do not vary with r_1 since both the true shares computed via $\eta_a(f)$ and the estimates computed via $\zeta_a(C)$ do not depend on the r_a . Importantly, the error is small despite the extreme disparity between popularity laws. The absolute error increases with the ratio T_1/T_2 but remains relatively small (e.g., $\text{err}_2 < 15\%$ for $T_1/T_2 < 10$). These results suggest it may be possible to design a practical subsidy scheme that combines the optimality of (9) and (10) with the verifiability of (8) and (13). Such a design is beyond present scope, however.

V. OPTIMIZING THE TRADEOFF

We first formulate the general tradeoff optimization problem under capacity constraints before elaborating a distributed solution based on Lagrangian relaxation for the particular network instance of Fig. 1. Numerical results demonstrate how the proposed method works on a toy network example.

A. General problem formulation

We assume content placement is optimized on a daily basis. CPs have detailed demand forecasts while ANOs are aware of capacity limits. Externally fixed costs are known to both CPs and ANOs. Optimal placements are realized at the start of the day in an off-peak traffic period. Charges and subsidies are calculated at the end of the day based on allocated cache capacities, measured peak period demand and respective unit resource prices.

We use the notation introduced in Sec. II-C with the addition of storage and bandwidth capacity limits for each node $n \in \mathcal{N}$, denoted S_n and B_n , respectively. We effectively assimilate the path from n to its parent to a link of capacity B_n . CPs, $k \in \mathcal{K}$, and ANOs, $a \in \mathcal{A}$, are implicitly identified by disjoint subsets of contents, $f \in \mathcal{F}_k$, and leaf nodes, $l \in \mathcal{L}_a$, respectively. All intermediate nodes, $\mathcal{I} = \mathcal{N} \setminus \{\mathcal{L}, 0\}$, also belong exclusively to one ANO. Let \mathcal{I}_a be the set of intermediate nodes belonging to ANO a and denote all nodes belonging to a by $\mathcal{N}_a = \mathcal{L}_a \cup \mathcal{I}_a$.

We also need notation for the cost of a path from each node n to the source: $b_{n+} = \sum_{n \leq m \leq 0} b_m$ where $n \leq m \leq 0$ covers the set of links on the path from n to the source. Costs are interpreted here as usage based *prices*. Each ANO can fix its own prices but, for the sake of neutrality, these are applied equally to all CPs.

Content placement is defined by indicator variables: $x_n^f = 1$ if content f is stored in node n and $x_n^f = 0$ otherwise. We additionally define delivery variables: $y_{ln}^f = 1$ if requests from leaf l for content f are delivered from node n and $y_{ln}^f = 0$ otherwise. In the optimal answer, we want to have $y_{ln}^f = 1$ if n is the first node such that $x_n^f = 1$ on the path from l to the source. The utility of a placement is equal to the overall savings relative to a network without caches. We seek to solve the following problem.

$$\begin{aligned} (\text{ILP}) \quad & \text{maximize } U(x, y) = \sum_{f \in \mathcal{F}} U(x^f, y^f), \text{ where} \\ & U(x^f, y^f) = \sum_{l \in \mathcal{L}} \sum_{l \leq n \leq 0} \lambda_l^f y_{ln}^f b_{n+} - \sum_{n \in \mathcal{N}} x_n^f s_n, \end{aligned} \quad (15)$$

subject to placement and delivery constraints for each content $f \in \mathcal{F}$,

$$\sum_{n \in \mathcal{N}} y_{ln}^f \leq 1, \forall l, \quad (16)$$

$$y_{ln}^f \leq x_n^f, \forall l, \forall n, \quad (17)$$

$$x_n^f \in \{0, 1\}, \forall n, \quad (18)$$

$$y_{ln}^f \in \{0, 1\}, \forall l, n, \quad (19)$$

and capacity constraints,

$$\sum_f x_n^f \leq S_n, \forall n, \quad (20)$$

$$\sum_f \lambda_l^f (1 - \sum_{l \leq m \leq n} y_{lm}^f) \leq B_n, \forall n. \quad (21)$$

This integer linear program is a generalization of the NP-hard problem of optimal content placement in cache hierarchies [21] (where only storage constraints (20) would be imposed). Before presenting a distributed solution based on Lagrangian relaxation, we discuss the constraints that apply in our particular instance.

B. A specific instance

We consider the three-tier network of Fig. 1. The root node 0 is the CO and, as assumed in Sec. IV, storage and bandwidth are unlimited with externally fixed prices s_0 and b_0 , respectively. Intermediate nodes, $i \in \mathcal{I}$, are also assumed able to elastically provide as much storage as needed at fixed unit price s_i . On the other hand, we suppose storage in leaf nodes is limited to $S_l < \infty$, for $l \in \mathcal{L}$, typically for technological reasons.

Bandwidth in the access network is supposed to be provisioned based on demand forecasts with a lead time of several months. This means bandwidth is limited but typically more than might be optimal at any considered instant within the planning cycle since demand has an increasing trend. We assume network planners are competent and ensure that no bandwidth constraint is violated when available cache space is optimally used, *i.e.*, the problem does have a feasible solution.

For capacity limited resources, the ANOs fix prices that induce CPs to optimally share available capacity. One objective would be, for instance, to fully use the limited cache in a wireless base station. Similarly, the bandwidth price on the link from an intermediate node i to the CO can be set to ensure maximum utilization and thus minimize the cost of storage at that node. Storage and bandwidth are effectively *substitutable* resources: the relative elasticity of storage provision at the intermediate nodes enables more efficient use of provisioned bandwidth.

C. A distributed solution

The complexity of (ILP) and the separate content placement and price setting roles of the CPs and ANOs, respectively, impose a distributed solution. This is possible through dual Lagrangian decomposition by relaxing constraints (20) and (21) and applying the sub-gradient method to derive an iterative procedure. The relaxed problem is then a set of UFL instances

as introduced in Sec. II-C and can be solved independently by the CPs.

First introduce notation for the cache size used by CP k at node n ,

$$C_{nk} = \sum_{f \in \mathcal{F}_k} x_n^f, \quad (22)$$

and the residual amount of CP k demand routed over link n ,

$$\Lambda_n^k = \sum_{l \in \mathcal{L}} \sum_{f \in \mathcal{F}_k} \lambda_l^f (1 - \sum_{l \leq m \leq n} y_{lm}^f). \quad (23)$$

For link 0, from CO to the source, we need the per-ANO partition of Λ_0^k . Let

$$\Lambda_{0a}^k = \sum_{l \in \mathcal{L}_a} \sum_{f \in \mathcal{F}_k} \lambda_l^f (1 - \sum_{l \leq m \leq 0} y_{lm}^f). \quad (24)$$

Finally, let ζ_{ak} be the ANO a share of the CO cache cost of CP k , computed from (9) and (10),

$$\zeta_{ak} = \frac{1}{C_{0k}} \sum_{f \in \mathcal{C}_{0k}} \frac{\sum_{l \in \mathcal{L}_a} \lambda_l^f y_{l0}^f}{\sum_{l \in \mathcal{L}} \lambda_l^f y_{l0}^f}. \quad (25)$$

For our specific network instance, utility can then be expressed as,

$$U(x, y) = \sum_{a \in \mathcal{A}} U_a(x, y), \quad (26)$$

where

$$U_a(x, y) = \sum_{k \in \mathcal{K}} \left(\left(\sum_{l \in \mathcal{L}_a} T_l^k - \Lambda_{0a}^k \right) b_0 - \zeta_{ak} C_{0k} s_0 - \sum_{i \in \mathcal{I}_a} C_{ik} s_i \right), \quad (27)$$

and $T_l^k = \sum_{f \in \mathcal{F}_k} \lambda_l^f$ is the demand for CP k at leaf l . The only costs appearing in (26) are b_0 , s_0 and s_i for $i \in \mathcal{I}$. Capacity limited resources are considered as sunk costs and the corresponding b_n and s_n are set to zero. Their utilization is controlled by shadow prices in the form of Lagrange multipliers.

Introduce the Lagrangian multipliers $\sigma_l \geq 0$ and $\beta_n \geq 0$ for the relaxed constraints (20) and (21), respectively, and define the Lagrangian,

$$L(x, y, \beta, \sigma) = \sum_{a \in \mathcal{A}} L_a(x, y, \beta, \sigma),$$

where

$$L_a(x, y, \beta, \sigma) = U_a(x, y) - \sum_{l \in \mathcal{L}_a} \sigma_l \left(\sum_{k \in \mathcal{K}} C_{lk} - S_l \right) - \sum_{n \in \mathcal{N}_a} \beta_n \left(\sum_{k \in \mathcal{K}} \Lambda_n^k - B_n \right), \quad (28)$$

and $U_a(x, y)$ is given by (27). The dual of (ILP) can then be written,

$$(\mathbf{DP}) \quad \min_{\beta \geq 0, \sigma \geq 0} \max_{x, y} L(x, y, \beta, \sigma), \quad (29)$$

where x and y are subject to constraints (16) to (19). We propose a distributed solution for (DP) where each CP places content to maximize L for given values of β and σ while the

Algorithm 1 Online algorithm to optimize utility (26). O denotes the orchestrator.

Result: Optimal x_n^f and y_{ln}^f , $\forall f, n, l$

Initialization:

O: sets $\tau = 1$, $\text{LB} = 0$, $\text{UB} = \sum_k \sum_l T_l^k b_0$, ϵ and γ .

ANO a : chooses initial $\sigma_l(0) \geq 0$, $\beta_n(0) \geq 0$.

- 1: **while** $\text{UB} - \text{LB} > \epsilon$ and $\tau < \tau_{\max}$ **do**
- 2: CP k : calculates optimal $x^f(\tau)$ and $y^f(\tau)$, $\forall f \in \mathcal{F}_k$
- 3: communicates the results (22), (23), (24) and (25) to each ANO a and to O.
- 4: O: $\text{LB} \leftarrow \max\{\text{LB}, U(x(\tau), y(\tau))\}$.
- 5: $\text{UB} \leftarrow \min\{\text{UB}, L(x(\tau), y(\tau), \beta(\tau), \sigma(\tau))\}$.
- 6: updates step size using (30) and sends to ANOs.
- 7: ANO a : updates shadow prices using (33).
- 8: $\tau = \tau + 1$

ANOs iteratively adjust these values using the sub-gradient method [22]. More precisely, we suppose the CPs and ANOs perform the distributed optimization at the start of each day, exchanging information in rounds of primal and dual update cycles. A third party ‘orchestrator’ (an algorithm executed in the CO, say) is additionally informed of iteration outcomes to determine step sizes and stopping conditions. The respective operations are summarized in Algorithm 1 and detailed below.

1) *Primal updates:* The CPs have private demand estimates λ_l^f , know fixed prices b and s and receive proposed shadow prices β and σ from the respective ANOs. Using these data, CP $k \in \mathcal{K}$ computes the optimal placement of each content $f \in \mathcal{F}_k$ by solving the UFL problem of Sec. II-C.

After performing its part of iteration τ , $\tau \geq 1$, each CP k , communicates the UFL results to each ANO a and the orchestrator in the form of cache capacities $C_{nk}(\tau)$ and residual traffic demands $\Lambda_n^k(\tau)$ for $n \in \mathcal{N}_a$, CO cache capacity $C_{0k}(\tau)$ and ANO cost share $\zeta_{ak}(\tau)$, and residual transit traffic $\Lambda_{0a}^k(\tau)$.

2) *Orchestration:* The orchestrator is aware of capacity constraints and verifies the feasibility of the CP results with regard to the primal problem (ILP). If feasible, these constitute a possible solution for the (ILP). The objective functions L and U are evaluated and compared to the current upper and lower bounds, UB and LB , respectively. As the bounds are updated, the CPs and ANOs must be informed to retain their corresponding decision variables.

The orchestrator stops iterations if the found solution has an absolute error of less than ϵ or a limit number of iterations has been attained. In either case the ANOs and CPs are informed and can implement the currently best feasible LB solution. If not, the orchestrator calculates a step size $\delta(\tau)$ for the next iteration using the Polyak formula [22] and sends it to all ANOs,

$$\delta(\tau) = \gamma \frac{|L(x(\tau), y(\tau), \beta(\tau), \sigma(\tau)) - \text{LB}|}{\|\nabla(\tau)\|^2}, \quad (30)$$

where $\gamma \geq 0$ is a scale factor and $\|\nabla(\tau)\|^2 = \sum_{n \in \mathcal{N}} \nabla_n^B(\tau)^2 + \sum_{l \in \mathcal{L}} \nabla_l^S(\tau)^2$. $\nabla_n^B(\tau)$ and $\nabla_l^S(\tau)$ are the sub-gradients at itera-

tion τ ,

$$\nabla_n^B(\tau) = \sum_{k \in \mathcal{K}} \Lambda_n^k(\tau) - B_n, \text{ for } n \in \mathcal{N} \setminus 0, \quad (31)$$

$$\nabla_l^S(\tau) = \sum_{k \in \mathcal{K}} C_{lk}(\tau) - S_l, \text{ for } l \in \mathcal{L}. \quad (32)$$

3) *Dual updates:* Each ANO a adjusts the shadow prices β_n and σ_n for capacity limited resources $n \in \mathcal{N}_a$ to more closely match the capacity constraints. Specifically, ANO a computes the sub-gradients $\nabla_n^B(\tau)$ for $n \in \mathcal{N}_a$ and $\nabla_l^S(\tau)$ for $l \in \mathcal{L}_a$. Using step size (30) communicated by the orchestrator, it computes

$$\begin{aligned} \beta_n(\tau + 1) &= [\beta_n(\tau) + \delta(\tau) \nabla_n^B(\tau)]^+, \text{ for } n \in \mathcal{N}_a, \\ \sigma_l(\tau + 1) &= [\sigma_l(\tau) + \delta(\tau) \nabla_l^S(\tau)]^+, \text{ for } l \in \mathcal{L}_a, \end{aligned} \quad (33)$$

where $[x]^+$ denotes the maximum of x and 0, and sends these new values to the CPs and the orchestrator.

Convergence to the optimal dual solution (DP) using the Polyak step size rule has been empirically demonstrated [22] but may take many iterations for large problems. Any feasible solution after a certain number of iterations is likely to be satisfactory, however, since capacity constraints are not usually tight (e.g., bandwidth is also used for other traffic) and slight under-utilization does not have serious consequences.

Due to the integer nature of decision variables, it is sometimes non-trivial to find feasible solutions. However, it often happens that the primal solution will be nearly feasible. In these cases, the orchestrator can project the solution to a fair, sub-optimal feasible solution for any violated constraint and require the CP to update the other related cache size and traffic demand variables. Alternative more sophisticated projection heuristics, typically requiring a greater degree of cooperation between orchestrator and CPs, might also be applied [22].

D. Settlements

The computed cache capacities C_{nk} are reserved for the day. Demands T_l^k and Λ_n^k are only estimates, however, and ANO charges and CP subsidies are based on measured busy period demand denoted \tilde{T}_l^k and $\tilde{\Lambda}_n^k$. The ANOs pay for storage in intermediate nodes $i \in \mathcal{I}_a$ and the CO, and for transit bandwidth at given rates s_i , s_0 and b_0 , respectively. They also pay a subsidy to the CPs equal to a fraction of their cost savings calculated using these fixed rates, the shadow prices and the measured traffic. ANO a pays CP k the following subsidy,

$$\begin{aligned} \text{sub}_{ak} &= r_{ak} \times \left(\sum_{l \in \mathcal{L}_a} \tilde{T}_l^k (\beta_l + \beta_{p(l)} + b_0) - \sum_{n \in \mathcal{N}_a} (\tilde{\Lambda}_n^k \beta_n) - \right. \\ &\quad \left. - \tilde{\Lambda}_{0a}^k b_0 - \sum_{l \in \mathcal{L}_a} (C_{lk} \sigma_l) - \sum_{i \in \mathcal{I}_a} (C_{ik} s_i) - \zeta_{ak} C_{0k} s_0 \right), \end{aligned} \quad (34)$$

where $p(l) \in \mathcal{I}_a$ is the parent of leaf l and r_{ak} is the proportionate share of overall savings negotiated between ANO a and CP k (see Sec. IV).




Fig. 5: Algorithm 1 applied to a symmetric network with 2 ANOs and 2 CPs; $e_1 = 100, e_2 = 10, b_0 = \4 per Mb/s per month, $s_0 = s_i = \$0.03$ per GB per month.

E. Numerical application

We consider a toy network example to illustrate an application of the method. As in Sec. II-D, to simplify the presentation of results, we consider a symmetric 3-tier access network where all leaves have the same demand and all links and nodes at the same level have the same capacity. Two ANOs have dedicated leaf and intermediate node infrastructure and access the Internet via a common CO hub. Both ANOs have identical demand and network capacities. The fanouts of each ANO are $e_1 = 100$ and $e_2 = 10$, as in the example of Fig. 3a.

ANO users access the content of two CPs both of which have a distinct content catalogue of 10^7 one megabyte files with a Zipf(0.8) popularity distribution. Peak demand for CP 1 content from the users of each ANO is 10 Gb/s while demand for CP 2 is twice as much.

The fixed cost of storage at the CO, s_0 , and at the intermediate nodes, s_i , is \$0.03 per GB while the cost of transit bandwidth b_0 is set to \$4 per Mb/s. These are monthly rates as in Sec. II-D. The other resources are provisioned by the ANOs and are considered as sunk costs. We derive optimal utilities, using Algorithm 1, as functions of the capacities of the other resources, namely leaf storage S_l , leaf bandwidth B_l and intermediate link bandwidth B_i . The negotiated fractional shares r_{ak} of ANO savings used to subsidize the CPs are arbitrarily set to 0.5 in all cases.

Fig. 5a shows how the intermediate link shadow price β_i varies as a function of capacity B_i while leaf storage and bandwidth are fixed. The figure shows that β_i decreases as available capacity increases, going to zero when link capacity ceases to be constraining. Shadow prices correspond as usual to marginal utilities and reflect the quality of the ANO dimensioning. They would coincide to real costs if the network were optimally sized.

Fig. 5b plots realized gains as a function of B_i for fixed S_l and B_l and Fig. 5c plots the same values as a function of S_l for fixed B_l and B_i . The figures show total utility (26), the subsidies received by each CP from (34), and the remaining ANO utility after deduction of the subsidy (identical for both ANOs). Total utility increases slightly as capacity grows but is not highly sensitive due to the low cost of storage, especially at the shared CO. Higher capacity leads to smaller shadow

prices and a greater share of the gains for the ANOs. The shares do not reflect the 0.5 ratio r_{ak} when capacities are low since shadow prices then contribute significantly to subsidies but not to overall utility. CP 2 has higher demand than CP 1 and therefore yields greater cost savings. This is why its subsidy is higher.

VI. RELATED WORK

The timely survey by Paschos *et al.* on the role of caching in networks usefully highlights the significant infrastructure economies realizable by trading off cache memory for access network bandwidth [23]. Cited works by Borst *et al.* [24] and Poularakis *et al.* [21] on optimal content placement in cache hierarchies are particularly relevant. However, this prior work does not explain how the optimal tradeoff can be realized when CPs, who have exclusive knowledge of demand, are hardly motivated to use this to reduce the costs incurred by the network operator.

The business relation between ISPs and CPs is well-known to be problematic and has given rise to much research and discussion, notably on the issue of network neutrality. The 2009 analysis of the content delivery two-sided market by Musacchio *et al.* [25] is still relevant today while forthcoming network “cloudification” does not appear to bring obvious simplifications (e.g., Tang and Ma [26] and Hu *et al.* [27]). In our work we suppose content placement takes place under the presently dominant, one-sided pricing model where end-user charges pay for access network costs. CPs therefore have little natural incentive to cooperate in minimizing these costs.

There is relatively little work that seeks to estimate the quantitative value of the memory bandwidth tradeoff. Kelly and Reeves [28] and Cidon *et al.* [29] made early contributions brought up to date by Erman *et al.* [30] and Roberts and Sbihi [5]. The present work extends the analysis of Elayoubi and Roberts [2] for an isolated cache by considering a hierarchy of cache locations. More significantly, we additionally design pricing and value sharing mechanisms that incite network operators and CPs to cooperatively realize the optimal tradeoff.

Game theory has been widely used to analyse the economic relation between ISPs and CPs. Much of this literature makes the assumption that CPs have an incentive to cache contents

close to end-users in the form of improved QoE, either through reduced latency or enhanced throughput. This is the case of recent work by Gourdin *et al.* [31], Mitra *et al.* [32], Mitra and Sridhar [33] and Douros *et al.* [19], for example. We disagree that QoE is a sufficiently discriminating criterion since latency in the access network is hardly impacted by negligible differences in propagation time while throughput and storage access times can and should be controlled by adequate provisioning. We deduce the need to introduce subsidies, in the form of side-payments from network operators to CPs, in order to realize the significant savings brought by optimizing the tradeoff. Game theory is used to determine how the savings should be shared between CPs and network operators.

When storage capacity or network bandwidth is limited, it is necessary to partition the resource between multiple CPs. Optimal partitioning of a limited capacity cache was considered by Dehghan *et al.* [34] and Araldo *et al.* [35]. Both papers propose iterative schemes where partitions are adjusted by the network operator based on observed per-CP performance. In our network model, optimal cache partitioning is realized by iteratively adjusting the price of storage relative to the price of bandwidth. The CP subsidy is such that they have an incentive to store the overall most popular contents, up to the capacity limit. Dynamic pricing for bandwidth sharing is the basis of network utility maximization (NUM) as introduced by Kelly *et al.* [36] and applied, in particular, to ISP-CP interaction by Hande *et al.* [37]. NUM is realized through a Lagrangian decomposition of the optimization into interacting user and network problems. In our proposal the tradeoff optimization is split between CPs placing content to optimize their subsidy while network operators fix prices to maximize utilization of their limited capacity pre-installed resources.

VII. CONCLUSION

We have identified the need for network operators to financially reward CPs for placing content in access network caches thus realizing an advantageous memory for bandwidth tradeoff. The considerable potential savings from an optimal tradeoff are currently not realized since the CPs, who exclusively possess the necessary detailed demand data, have no natural incentive to use this to make the optimal placement. In our proposal, each ANO would give a subsidy to each CP that is proportional to the realized savings. The actual proportion, between 0 and 100%, would be determined bilaterally.

We designed a value sharing scheme where CPs maximize their subsidy by optimally placing content items in an access network cache hierarchy rooted at the central office. We determined a value distribution that specifies how the cost of the CO cache and corresponding CP subsidy should be divided among multiple ANOs. This distribution is independent of the individually negotiated proportions of savings handed over to the CP.

We proposed a distributed iterative approach to optimize the tradeoff based on Lagrangian decomposition. CPs compute their optimal content placements in the cache hierarchy, given unit bandwidth and storage prices, while ANOs fix shadow

prices to maximize utilization of limited capacity resources. Neutrality is assured since prices are applied uniformly and content items are placed optimally independently of the CP to which they belong.

The present cache subsidy proposal is novel and considerable scope to extend this preliminary analysis remains. Our simple network model excludes certain aspects of existing and future access networks that impact the tradeoff. For instance, overlapping base station coverage areas or a topology with cross links would allow cooperative caching [17]. Infrastructure sharing between ANOs in 5G networks would imply further cost and subsidy partitions [16].

Our analysis and algorithm design can certainly be improved in several directions. We have not thoroughly evaluated the convergence speed and optimality gap of the proposed distributed optimisation. The subsidy has been designed so that CPs maximize their gain by making optimal placements but it remains to fully evaluate scope for gaming the system by either ANOs or CPs. Fallback actions should be defined when some ANO or CP does not participate or behaves irrationally.

Last but not least, a practical cache subsidy scheme needs to be acceptable to ANOs and CPs. Network operators already complain that CPs do not pay them sufficiently and will not at all like the idea of payments going in the opposite direction. CPs, on the other hand, are quite happy with their lucrative business models and may not see any pressing need to change. However, optimizing the network infrastructure is a worthwhile societal goal leading ultimately to lower charges to end users and greater efficiency for both network operators and CPs. Cache subsidies are a viable means to achieve this goal.

ACKNOWLEDGMENT

The work presented in this article has benefited from the support of NewNet@Paris, Cisco's Chair "NETWORKS FOR THE FUTURE" at Telecom ParisTech (<http://newnet.telecom-paristech.fr>). Any opinions, findings or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of partners of the Chair.

REFERENCES

- [1] L. Peterson, A. Al-Shabibi, T. Anshutz, S. Baker, A. Bavier, S. Das, J. Hart, G. Palukar, and W. Snow, "Central office re-architected as a data center," *IEEE Commun. Mag.*, vol. 54, no. 10, pp. 96–101, 2016.
- [2] S.-E. Elayoubi and J. Roberts, "Performance and cost effectiveness of caching in mobile access networks," in *Proc. 2nd ACM Conf. on Information-Centric Netw. (ICN)*, 2015, pp. 79–88.
- [3] C. Imbrenda, L. Muscariello, and D. Rossi, "Analyzing cacheable traffic in ISP access networks for micro CDN applications via content-centric networking," in *Proc. 1st ACM Conf. on Information-Centric Netw. (ICN)*, 2014, pp. 57–66.
- [4] F. Olmos and B. Kauffmann, "An inverse problem approach for content popularity estimation," in *Proc. 9th EAI Int. Conf. on Performance Evaluation Methodologies and Tools*, ser. VALUETOOLS'15. ICST, 2016, pp. 33–40.
- [5] J. Roberts and N. Sbihi, "Exploring the memory-bandwidth tradeoff in an information-centric network," in *Proc. 25th Int. Teletraffic Congress (ITC)*. IEEE, 2013, pp. 1–9.
- [6] "Google Cloud," <https://cloud.google.com>, accessed 2019-08-15.
- [7] G. Cornuéjols, G. Nemhauser, and L. Wolsey, *The Uncapacitated Facility Location Problem*, ser. Discrete Location Theory. John Wiley and Sons Inc., 1991, pp. 119–171, in P Mirchandani and R Francis editors.

- [8] R. Shah and M. Farach-Colton, "Undiscretized dynamic programming: faster algorithms for facility location and related problems on trees," in *Proc. 13th annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 2002, pp. 108–115.
- [9] C. Perillo, A. Antonopoulos, and C. Verikoukis, "Panel dataset description for econometric analysis of the isp-ott relationship in the years 2008–2013," *arXiv preprint arXiv:1612.06451*, 2016.
- [10] "The global internet phenomena report," White Paper, Sandvine, October 2018.
- [11] J.-C. Rochet and J. Tirole, "Two-sided markets: a progress report," *The RAND J. of Economics*, vol. 37, no. 3, pp. 645–667, 2006.
- [12] R. T. Ma, D. M. Chiu, J. C. Lui, V. Misra, and D. Rubenstein, "On cooperative settlement between content, transit, and eyeball internet service providers," *IEEE/ACM Trans. Netw.*, vol. 19, no. 3, pp. 802–815, 2011.
- [13] T. Böttger, F. Cuadrado, G. Tyson, I. Castro, and S. Uhlig, "Open connect everywhere: A glimpse at the internet ecosystem through the lens of the netflix cdn," *ACM SIGCOMM Computer Communication Review*, vol. 48, no. 1, pp. 28–34, 2018.
- [14] K.-K. Yap, M. Motiwala, J. Rahe, S. Padgett, M. Holliman, G. Baldus, M. Hines, T. Kim, A. Narayanan, A. Jain *et al.*, "Taking the edge off with espresso: Scale, reliability and programmability for global internet peering," in *Proc. Conf. of the ACM Special Interest Group on Data Communication*, 2017, pp. 432–445.
- [15] F. Wohlfart, N. Chatzis, C. Dabanoglu, G. Carle, and W. Willinger, "Leveraging interconnections for performance: the serving infrastructure of a large cdn," in *Proc. Conf. of the ACM Special Interest Group on Data Communication*, 2018, pp. 206–220.
- [16] J. G. Andrews, S. Buzzi, W. Choi, S. V. Hanly, A. Lozano, A. C. K. Soong, and J. C. Zhang, "What will 5g be?" *IEEE J. Sel. Areas Commun.*, vol. 32, no. 6, pp. 1065–1082, 2014.
- [17] J. Krolkowski, A. Giovanidis, and M. Di Renzo, "A decomposition framework for optimal edge-cache leasing," *IEEE J. Sel. Areas Commun.*, vol. 36, no. 6, pp. 1345–1359, 2018.
- [18] W. Saad, Z. Han, M. Debbah, A. Hjørungnes, and T. Basar, "Coalitional game theory for communication networks," *IEEE Signal Process. Mag.*, vol. 26, no. 5, pp. 77–97, 2009.
- [19] V. G. Douros, S. E. Elayoubi, E. Altman, and Y. Hayel, "Caching games between content providers and internet service providers," *Performance Evaluation*, vol. 113, pp. 13–25, 2017.
- [20] E. Kalai and D. Samet, "On weighted shapley values," *Int. J. of Game Theory*, vol. 16, no. 3, pp. 205–222, 1987.
- [21] K. Poularakis and L. Tassiulas, "On the complexity of optimal content placement in hierarchical caching networks," *IEEE Trans. Commun.*, vol. 64, no. 5, pp. 2092–2103, 2016.
- [22] M. L. Fisher, "The lagrangian relaxation method for solving integer programming problems," *Management Science*, vol. 27, no. 1, pp. 1–18, 1981.
- [23] G. S. Paschos, G. Iosifidis, M. Tao, D. Towsley, and G. Caire, "The role of caching in future communication systems and networks," *IEEE J. Sel. Areas Commun.*, vol. 36, no. 6, pp. 1111–1125, 2018.
- [24] S. Borst, V. Gupta, and A. Walid, "Distributed caching algorithms for content distribution networks," in *Proc. IEEE INFOCOM*, 2010, pp. 1–9.
- [25] J. Musacchio, G. Schwartz, and J. Walrand, "A two-sided market analysis of provider investment incentives with an application to the net-neutrality issue," *Review of Network Economics*, vol. 8, pp. 22–39, 2009.
- [26] J. Tang and R. T. Ma, "Regulating monopolistic isps without neutrality," *IEEE J. Sel. Areas Commun.*, vol. 37, no. 7, pp. 1666–1680, 2019.
- [27] X. Hu, G. Kesidis, B. Heidarpour, and Z. Dziong, "Media delivery competition with edge cloud, remote cloud and networking," *Network Games, Control, and Optimization*, pp. 71–87, 2019.
- [28] T. Kelly and D. Reeves, "Optimal web cache sizing: Scalable methods for exact solutions," *Computer Communications*, vol. 24, no. 2, pp. 163–173, 2001.
- [29] I. Cidon, S. Kutten, and R. Soffer, "Optimal allocation of electronic content," *Computer Networks*, vol. 40, no. 2, pp. 205–218, 2002.
- [30] J. Erman, A. Gerber, M. Hajiaghayi, D. Pei, S. Sen, and O. Spatscheck, "To cache or not to cache: The 3g case," *IEEE Internet Comput.*, vol. 15, no. 2, pp. 27–34, 2011.
- [31] E. Gourdin, P. Maillé, G. Simon, and B. Tuffin, "The economics of cdns and their impact on service fairness," *IEEE Trans. Network Service Management*, vol. 14, no. 1, pp. 22–33, 2017.
- [32] D. Mitra, Q. Wang, and A. Hong, "Emerging internet content and service providers' relationships: Models and analyses of engineering, business and policy impact," in *Proc. IEEE INFOCOM*, 2017, pp. 1–9.
- [33] D. Mitra and A. Sridhar, "The case for formation of isp-content providers consortiums by nash bargaining for internet content delivery," in *Proc. IEEE INFOCOM*, 2019.
- [34] M. Dehghan, W. Chu, P. Nain, D. Towsley, and Z.-L. Zhang, "Sharing cache resources among content providers: A utility-based approach," *IEEE/ACM Trans. Netw.*, vol. 27, no. 2, pp. 477–490, 2019.
- [35] A. Araldo, G. Dán, and D. Rossi, "Caching encrypted content via stochastic cache partitioning," *IEEE/ACM Trans. Netw.*, vol. 26, no. 1, pp. 548–561, 2018.
- [36] F. P. Kelly, A. K. Maulloo, and D. K. Tan, "Rate control for communication networks: shadow prices, proportional fairness and stability," *J. of the Operational Research society*, vol. 49, no. 3, pp. 237–252, 1998.
- [37] P. Hande, M. Chiang, R. Calderbank, and S. Rangan, "Network pricing and rate allocation with content provider participation," in *Proc. IEEE INFOCOM*, 2009, pp. 990–998.