



HAL
open science

Optimal Cache Leasing from a Mobile Network Operator to a Content Provider

Jonatan Krolikowski, Anastasios Giovanidis, Marco Di Renzo

► **To cite this version:**

Jonatan Krolikowski, Anastasios Giovanidis, Marco Di Renzo. Optimal Cache Leasing from a Mobile Network Operator to a Content Provider. IEEE International Conference on Computer Communications (INFOCOM), Apr 2018, Honolulu, United States. 10.1109/infocom.2018.8486267. hal-01663134

HAL Id: hal-01663134

<https://hal.science/hal-01663134>

Submitted on 15 Jul 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Optimal Cache Leasing from a Mobile Network Operator to a Content Provider

Jonatan Krolikowski*, Anastasios Giovanidis[†] and Marco Di Renzo*

*CNRS-L2S, CentraleSupélec, Université Paris Sud, Université Paris-Saclay

[†]CNRS-LIP6, Université Pierre et Marie Curie, Sorbonne Universités, Paris, France

Email: jonatan.krolikowski@telecom-paristech.fr, anastasios.giovanidis@lip6.fr, marco.direnzo@lss.supelec.fr

Abstract—Caching popular content at the wireless edge is recently proposed as a means to reduce congestion at the backbone of cellular networks. The two main actors involved are Mobile Network Operators (MNOs) and Content Providers (CPs). In this work, we consider the following arrangement: an MNO pre-installs memory on its wireless equipment (e.g. Base Stations) and invites a unique CP to use them, with monetary cost. The CP will lease memory space and place its content; the MNO will associate network users to stations. For a given association policy, the MNO may help (or not) the CP to offload traffic, depending on whether the association takes into account content placement.

We formulate an optimization problem from the CP perspective, which aims at maximizing traffic offloading with minimum leasing costs. This is a joint optimization problem that can include any association policy, and can also derive the optimal one. We present a general exact solution using Benders decomposition. It iteratively updates decisions of the two actors separately and converges to the global optimum. We illustrate the optimal CP leasing/placement strategy and hit probability gains under different association policies. Performance is maximised when the MNO association follows CP actions.

I. INTRODUCTION

By 2020, wireless data traffic is estimated to reach roughly the 8-fold of its volume of 2015 [1]. Such increase is a challenge to mobile network operators (MNOs) as well as to the content providers (CPs). The MNOs and the CPs have different strategies to cope with these high demands: The MNOs, on the one hand, try to satisfy the increase in data demand by densifying the network with new tiers and by allowing cooperation among stations. However, this increase of wireless traffic can pose new problems to the wireless backhaul that are related to congestion. On the other hand, the CPs are on the receiving end of the content requests. Large CPs such as Youtube or Netflix store their data in huge data centers. For such a CP, a steep increase in data demand can be handled by massive infrastructure investments, i.e. upgrade of the data center capacity as well as installation of higher capacity data links to the surrounding network.

A recently studied alternative is to equip wireless nodes with caches (see [2]–[6]). The main purpose of cache memory installation is to ease backhaul traffic and its processing at the data centers by handling content requests from intermediate caches placed at the edge or inside the core network. This way, the user Quality of Service (QoS) can be improved because

cached content is downloaded with less delay from caches closer to the user.

Caching can be both of interest to the MNOs and the CPs. We consider a scenario in which an MNO has constructed and physically maintains caching infrastructure at Cached Base Stations (CBSs). The MNO makes three strategic decisions:

MNO-1) How much memory to install at each CBS?

MNO-2) Which price to set for the leasing of one cache unit?

MNO-3) Which user association policy to pursue?

The standard user association policy routes each user to its closest station because of the strongest channel quality. More involved policies allow load-balancing of users when these are located in the overlap of coverage areas of two or more CBSs.

A central assumption we make in this work is that the MNO puts the management of cache content into the hands of the CP. There are two reasons for this. The first is that the CP usually transfers data to wireless users via secure connections, e.g. https [7], so that the intermediate MNO cannot recognize requests and serve them from local caches. The second is that content placement policies may depend on spatio-temporal popularity data that the CP can gather and use. All content related information is in the hands of the CP.

In what follows, we take the point of view of the CP, which reacts to the MNO's cache installation, pricing scheme, and user association policy. The CP has to take two types of decisions:

CP-1) How much cache space to lease from the MNO at each CBS?

CP-2) Which content to place into the respective leased caches?

The decisions are based on the estimated spatio-temporal popularity of the CP's content. They are taken for a time period in which the system parameters are considered relatively static. For example, updates can take place in off-peak hours [8]. The aim of the CP is to find a cache leasing and content placement strategy which optimally weighs the total cache leasing costs against the savings from a reduction of traffic at the data centers: For every object stored in a cache, a price for memory space and content delivery has to be paid. But, at the same time, all users that are associated to the cache and ask for the cached content generate savings for the CP by relieving data centers [9]. To better understand the problem, consider two examples. In each case, the MNO installs unlimited cache

memory at all CBSs and sets a price per memory unit. Then a user association policy is announced:

CLOSEST: The MNO associates each wireless user to the geographically closest CBS. Based on offline estimates of content popularity in each association area, the CP can calculate the potential savings per CBS when caching each content file. The optimal strategy for the CP is to cache exactly those files for which the savings exceed the corresponding cache leasing costs. In this case the MNO associates users independently of their requests and CP decisions. Note that, for zero costs and unit file size, it is optimal to place the locally Most Popular Content in each CBS.

OPT-h: A different MNO association policy allows for more user requests to be served by the caches. If users can be associated to any single wireless node among those close enough (in contrast to a predetermined association, e.g. CLOSEST) each user has potential access to the set of files cached at all covering stations. It is then beneficial for the CP to cache different sets of content in neighboring nodes. To maximize those benefits, the MNO should associate users to CBSs in such a way that content requests are matched to the cached content. A special case for such user association is used for FemtoCaching [2].

The contributions of this work are summarized as follows:

- We introduce a model whose elements include cache leasing at CBSs, content placement decisions and user association.
- We formulate a mixed integer Network Utility Maximization (NUM) problem that aims to maximize the CP's benefits given the user association policy. The leasing and caching decisions are discrete, taken by the CP. The MNO association decision variables are fractional.
- As solution technique, we introduce Benders decomposition of the NUM which converges to the global optimum. One of its main advantages, aside optimality, is that it allows the separation of the user association problem from the cache leasing and content placement problems in the solution process. This separation of the solution process into two subproblems solved iteratively by the two actors is reminiscent of the work by Kelly [10]. Our solution is completely original for solving NUM edge-caching problems with non-linear utilities.
- We provide extensive evaluation of the optimal leasing, placement and pricing strategies for linear and concave objective functions. We further show the benefits in hit probability, when the MNO makes optimal association decisions taking content placement into account, compared to content agnostic strategies, like the closest node association policy.

The remainder of this paper is organized as follows: In Section II, we survey relevant literature to our problem. The system model is developed in Section III where we also state the general mixed integer NUM problem. The general solution based on Benders decomposition is provided in Section IV. Special formulations of the problem with linear or separable concave objective functions as well as different association

policies (CLOSEST, OPT-h) are presented in Section V. The Section also includes a discussion on how to include wireless resource sharing in the problem formulation. Extensive numerical evaluation of these problems and analysis of the findings is given in Section VI. Finally, Section VII concludes our work.

II. RELATED LITERATURE

There is extensive literature on the advantages of caching in wireless networks. However, just few works treat the joint leasing, placement and routing problem, none of which finds the exact optimal solution until now.

Shanmugam et al. [2] place content into caches of small cells to minimize the network delay. Decisions are taken based on popularity data. The users are associated to any covering station without taking resource constraints of the wireless nodes into consideration. In [3], Błaszczyszyn and Giovanidis develop a content placement policy which maximizes the hit ratio. This general probabilistic solution is not tailored to specific network topologies. According to Baştuğ, Bennis and Debbah [4], the globally most popular files are stored in wireless caches. Users are associated to the closest base station, not knowing if the requested content is stored in the cache or not. Poularakis, Iosifidis, and Tassiulas maximize in [8] the hit ratio by means of integer optimization. They introduce a bandwidth constraint limiting the number of users that can be connected to each cellular station. A polynomial-time approximation scheme is also provided. Deghan et al. [11] develop an approximation algorithm for the joint content placement and user association problem minimizing network delay. Their model controls if users are routed through a cached or an uncached path. The users on a cached path are always associated to the closest cache storing the content. Naveen et al. [12] provide an optimal placement and user association scheme with fractional content placement. In [13], the authors solve a NUM user association problem assuming that content is already placed in caches.

Considering cache leasing, Poularakis et al. [5] propose a business model where residential internet users lease part of their wireless bandwidth and storage capacity to the MNO in exchange for financial reimbursement. A joint optimization of incentive, content placement and routing policies provides offloading of backhaul traffic to local caches. However, their Lagrangian based solution does not converge to the global optimum due to weak duality (see p. 143 in [14]) and other solution techniques are necessary to solve the problem optimally. The idea of sharing backbone cache memory among different CPs is introduced by Araldo, Dan and Rossi [7]. The partitioning of the caches remains under the control of the internet service provider, while the CPs are allowed to establish secure connections between caches and users. Another work on the topic by Douros et al. is [6]. Paschos et al. [15] discuss the roles of MNOs, CPs and users as well as practical aspects of wireless caching such as the compilation of popularity data and the limitations arising from transmission encryption.

Related to the solution approach, Bektaş et al. [16] use Benders decomposition for joint placement and routing problems that have linear utilities and binary variables in the context of Content Delivery Networks. Their work uses a similar technique as in our work. The main difference is that we treat a wireless network and we use Generalised Benders' decomposition for non-linear NUM problems. Our problem here is mixed integer with continuous association variables. The decomposition in our paper has a natural business interpretation.

III. PROBLEM STATEMENT AND SYSTEM MODEL

A. Problem statement

The objective of the CP is to lease cache memory at the CBSs and place content into it such that the relation of its expected savings to the leasing cost is optimal. The savings are given by the function $h(\cdot)$ that takes as input the user association vector $\mathbf{y}^\Pi(\mathbf{x})$ where \mathbf{x} is the content placement action and Π is the MNO's association policy. The leasing costs at each CBS m are the product of leased units z_m times the price per unit q_m that is set by the MNO. An additional fee for the appropriate user association and content delivery can be included. Formally, the CP seeks a feasible tuple of vectors $(\mathbf{x}, \mathbf{z}) \in \mathcal{X}$ that maximizes the objective function $h(\mathbf{y}^\Pi(\mathbf{x})) - \sum_{m \in \mathcal{M}} q_m z_m$.

The CP's Cache Leasing and Content Placement problem (CLCP) can be formulated as the Non-Linear Mixed-Integer Problem (NLMIP)

$$\begin{aligned} \text{(CLCP)} \quad & \max_{\substack{(\mathbf{x}, \mathbf{z}) \in \mathcal{X} \\ \mathbf{y} \in \mathcal{Y}^\Pi}} h(\mathbf{y}) - \sum_{m \in \mathcal{M}} q_m z_m \\ \text{s. t.} \quad & y_{m,s,f} \leq N_{s,f} x_{m,f}, \quad \forall m, s, f, \end{aligned}$$

where m is a CBS, s is a planar network region and f is a data file. All components of this problem will be formally presented and explained in this section.

B. Cache Leasing and Content Placement

We consider a cellular communications network with a finite set \mathcal{M} of CBSs. Each CBS m is equipped with k_m memory units of size b_{MU} (in MBytes, e.g. 1000) which the CP can lease. Denoting the decision variable of how many cache units to lease (CP-1) at m by $z_m \in \mathbb{Z}_{\geq 0}$, the bounded availability of memory gives the constraint set

$$z_m \leq k_m \quad \forall m \in \mathcal{M}. \quad (1)$$

The vector of the cache leasing variables is $\mathbf{z} = (z_m)_{m \in \mathcal{M}}$.

Having leased cache space at the CBSs, the CP places content from a finite object catalog \mathcal{F} into the caches (CP-2). The decision to store content f in the cache of m will set the variable $x_{m,f}$ to 1, otherwise $x_{m,f} = 0$. The vector of content placement variables is $\mathbf{x} = (x_{m,f})_{m \in \mathcal{M}, f \in \mathcal{F}}$. Each file $f \in \mathcal{F}$ has a given file size b_f (in MBytes), and all file-sizes

are known. The limited capacity of the leased cache space gives the second constraint set

$$\sum_{f \in \mathcal{F}} b_f x_{m,f} \leq b_{\text{MU}} z_m \quad \forall m \in \mathcal{M}. \quad (2)$$

For convenience, we define the set of feasible tuples of leasing and placement vectors as

$$\mathcal{X} := \{(\mathbf{x}, \mathbf{z}) \in \{0, 1\}^{|\mathcal{M}||\mathcal{F}|} \times \mathbb{Z}_{\geq 0}^{|\mathcal{M}|} \mid (1), (2)\}.$$

C. Wireless Environment and User Association

Coverage Cells: Our communications model is the following: Each CBS has a planar 2D coverage cell. Users covered by a CBS receive a radio signal strong enough to be potentially associated to it. Coverage cells may overlap, thus offering the users multiple options for service from covering CBSs. However, we do not allow simultaneous service by more than one station, i.e. cooperative service is not possible.

Network Regions: The network area is partitioned into a set of regions. All positions in each region are assumed to experience the same radio conditions with respect to fading and interference. Furthermore, the MNO has a user association policy Π that allows for users in region s to potentially be associated to any CBS in $\mathcal{M}(s) \subseteq \mathcal{M}$, $|\mathcal{M}(s)| \geq 1$. (a) With the traditional policy that associates users to the closest station, there can be several covering CBSs for region s , but $\mathcal{M}(s)$ is the set consisting of only the closest station. (b) For the OPT-h policy, to be defined later, $\mathcal{M}(s)$ is the set of covering CBSs. In general, for policy Π , the set of regions is denoted by \mathcal{S}^Π (see Figure 1 for an example).

Content Popularity: For each region $s \in \mathcal{S}^\Pi$ and each content $f \in \mathcal{F}$, the expected number of users in s requesting f is considered to be known, measured or estimated by some process we do not consider here. It is denoted by $N_{s,f}$. The content popularity vector is $\mathbf{N} = (N_{s,f})_{s \in \mathcal{S}^\Pi, f \in \mathcal{F}}$.

User Association Variables and Constraints: In order to make optimal decisions, the CP has two types of information at its disposal: the popularity vector \mathbf{N} and the MNO's user association policy Π . Knowledge of \mathbf{N} and Π allows the CP to take decisions based on the expected association of users with CBSs. In the context of this work, we are only interested in *cache-hit* traffic, i.e. the traffic of users who find their request cached at the CBS they are associated to. The association vector of *cache-hit* users to the CBSs is $\mathbf{y} = (y_{m,s,f})_{m \in \mathcal{M}, s \in \mathcal{S}^\Pi(m), f \in \mathcal{F}}$, where $y_{m,s,f}$ represents the expected user traffic from region s requesting content f and associated with CBS m . $\mathcal{S}^\Pi(m)$ is the subset of regions whose users can potentially be associated to m according to Π . The vector \mathbf{y} has fractional non-negative entries.

User assignment is unique in the sense that an association to two or more CBSs is not allowed. The total population $N_{s,f}$ is distributed among the CBSs $\mathcal{M}(s)$, and some of it is potentially not associated to any CBS at all. Thus,

$$\sum_{m \in \mathcal{M}(s)} y_{m,s,f} \leq N_{s,f}, \quad \forall s \in \mathcal{S}^\Pi, f \in \mathcal{F}. \quad (3)$$

This constraint allows for possible splitting of the population $N_{s,f}$ among the CBSs in $\mathcal{M}(s)$. The set of assignment vectors feasible to this constraint set is denoted by

$$\mathcal{Y}^\Pi := \left\{ \mathbf{y} \in \mathbb{R}_{\geq 0}^{\sum_{m \in \mathcal{M}} |\mathcal{S}^\Pi(m)| |\mathcal{F}|} \mid (3) \right\}.$$

Since we are only interested in cache-hit traffic, $y_{m,s,f}$ can only be nonzero if $x_{m,f} = 1$, i.e. if object f is cached in station m . Since no more than the total population requesting content f in s can be included in $y_{m,s,f}$, the following constraint set is valid:

$$y_{m,s,f} \leq N_{s,f} x_{m,f}, \quad \forall m \in \mathcal{M}, s \in \mathcal{S}^\Pi(m), f \in \mathcal{F}. \quad (4)$$

This constraint set is very important since it couples MNO association variables with CP cache placement decisions.

D. CP savings

The CP uses the general *savings function* $h(\cdot)$ to measure user association \mathbf{y} . This function represents the savings (in €) obtained when users are associated with caches that store the requested content, thus avoiding use of its data centers.

In this paper, we solve CLCP for any monotonously increasing, continuously differentiable and concave savings function. Two choices for $h(\cdot)$ are particularly of interest:

- i) In case that the CP is solely interested in maximizing the hit ratio, it can choose $h(\cdot)$ as a *linear* function.
- ii) Choosing $h(\cdot)$ as the sum of *strictly concave* functions (one function per CBS), the CP can include aspects such as soft resource requirements and load-balancing. Additional communication conditions (e.g. fading and interference) can also be included in h .

The discussion over particular choices of $h(\cdot)$, which result in problems with different objectives, is postponed to Section V.

E. MNO policy

The way the MNO associates users to CBSs determines the association vector \mathbf{y} , and depends on the MNO's user association policy Π . The policies considered in this work are:

- 1) CLOSEST: Association to the closest covering CBS.
- 2) OPT-h: Association maximizing the CP's savings function. Observe that (4) implies that cache-hits (vector \mathbf{y}) depend on the placement \mathbf{x} in all cases. The resulting association vector is denoted by $\mathbf{y}^\Pi(\mathbf{x})$. If $\Pi = \text{CLOSEST}$, the association entry $y_{m,s,f}$ is positive only for users that find their content cached. However, association actions do not depend on placement \mathbf{x} .

On the other hand, if $\Pi = \text{OPT-h}$, the MNO fully cooperates with the CP in the sense that it always adapts its association vector \mathbf{y} to the placement \mathbf{x} such that the CP's savings function h is maximized. This is achieved by splitting traffic among CBSs given multi-coverage.

For the two policies $\Pi = \text{OPT-h}$ and $\Pi = \text{CLOSEST}$, the set of association regions \mathcal{S}^Π is different as explained in III-C (User Association). We illustrate this difference in the example

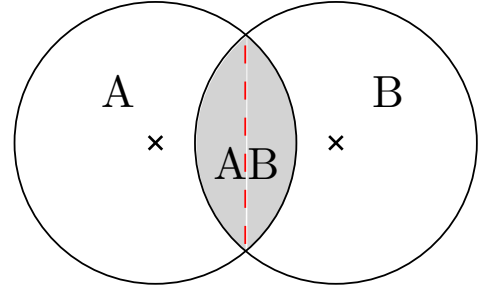


Fig. 1: In case $\Pi = \text{OPT-h}$, there are three regions A, B and AB. Users in region A and B can only be associated to their uniquely covering CBSs, respectively. Users in region AB can potentially be associated to any of the two CBSs. If $\Pi = \text{CLOSEST}$ (dashed line), there are two regions: A and the left part of AB contain traffic entirely associated to the left CBS, B and the right part of AB contain traffic belonging to the right CBS.

of Figure 1. For both policies, the association vector is the optimal solution to the User Association problem

$$\begin{aligned} \text{(UA-}\Pi\text{)} \quad \mathbf{y}^\Pi(\mathbf{x}) &= \arg \max_{\mathbf{y} \in \mathcal{Y}^\Pi} h(\mathbf{y}) \\ &\text{s.t.} \end{aligned} \quad (4).$$

This problem is convex, thus always tractable.

For CLOSEST, the association vector $\mathbf{y}^{\text{CLOSEST}}(\mathbf{x})$ follows immediately since $y_{m,s,f} = N_{s,f}$ if $x_{m,f} = 1$ and $y_{m,s,f} = 0$ otherwise. These values also are the optimum of UA-CLOSEST, since h is strictly increasing. For OPT-h, the solution of UA-OPT-h can be found by convex programming methods.

IV. SOLUTION

The high complexity of CLCP arises from the fact that it is a mixed-integer problem with non-linear objective. The solution technique that resolves this problem is Generalized Benders decomposition by Schrijver [17] and Geoffrion [18] which converges to the global optimum. As in the work of Elwalid, Mitra and Wang [19] who also use this technique in a different setting, we solve the non-linear part and the integer part separately. CLCP is a generalization of the Helper Decision Problem proved to be NP-complete in [2] even in the case of a linear savings function. Thus, our solution algorithm cannot be polynomial unless $P = NP$. However, a state-of-the-art MIP solver can be used for the solution. Its performance is shown in Section VI. In what follows, we give an overview over Benders decomposition applied to CLCP.

CLCP can be decomposed into two problems called Master and Slave. Master decides about cache leasing and content placement. Slave computes the optimal user association for a fixed content placement. We obtain

$$\text{(Master)} \quad \max_{(\mathbf{x}, \mathbf{z}) \in \mathcal{X}} h(\mathbf{y}(\mathbf{x})) - \sum_{m \in \mathcal{M}} q_m z_m,$$

where $h(\mathbf{y}(\mathbf{x}))$ is the objective value of

$$\begin{aligned} \text{(Slave)} \quad \mathbf{y}(\mathbf{x}) &= \arg \max_{\mathbf{y} \in \mathcal{Y}^\Pi} h(\mathbf{y}) \\ &\text{s.t.} \end{aligned} \quad (4).$$

Note that the Master problem can be treated by the CP, the Slave by the MNO. \mathcal{X} is discrete and finite and \mathcal{Y}^Π is compact and convex. Slave is the UA-II problem from Section III-E which is generally non-linear. Thus, Master cannot be solved directly. Generalized Benders decomposition deals with this problem by solving a sequence of Slave problems for different values of \mathbf{x} (and \mathbf{z}). The solutions to Slave are used to construct (linear) Benders cuts that define approximations to Slave, thus linearizing the problem.

Let $\{(\mathbf{x}^t, \mathbf{z}^t) \in \mathcal{X} \mid t = 1, \dots, T\}$ be a set of vector tuples feasible to Master for some $T \geq 0$. Let $\mathbf{y}^t := \mathbf{y}(\mathbf{x}^t)$ denote a corresponding vector that optimizes Slave for given \mathbf{x}^t , and $\boldsymbol{\lambda}^t = (\lambda_{m,s,f}^t)_{m \in \mathcal{M}, s \in \mathcal{S}^\Pi(m), f \in \mathcal{F}}$ be the vector of Lagrangian multipliers corresponding to the constraints (4). The existence of such vector \mathbf{y}^t follows from the fact that $N_{s,f} \geq 0$ holds true for all $s \in \mathcal{S}^\Pi$, $f \in \mathcal{F}$. The computability of \mathbf{y}^t and $\boldsymbol{\lambda}^t$ are assumed here, and discussed for special cases in Section V.

Slave is a convex problem. Thus, the duality theorem of convex programming implies

$$h(\mathbf{y}(\mathbf{x})) \leq h(\mathbf{y}^t) + \sum_{m \in \mathcal{M}} \sum_{s \in \mathcal{S}^\Pi(m)} \sum_{f \in \mathcal{F}} \lambda_{m,s,f}^t (N_{s,f} x_{m,f} - y_{m,s,f}^t)$$

for all feasible vectors \mathbf{x} . This upper bound to Slave is called *Benders cut*. Reformulated, we get

$$h(\mathbf{y}(\mathbf{x})) \leq \Gamma^t + (\mathbf{v}^t)' \mathbf{x}, \quad (5)$$

where

$$\Gamma^t := h(\mathbf{y}^t) - \sum_{m \in \mathcal{M}} \sum_{s \in \mathcal{S}^\Pi(m)} \sum_{f \in \mathcal{F}} \lambda_{m,s,f}^t y_{m,s,f}^t$$

and $(\mathbf{v}^t)'$ is the transpose of $\mathbf{v}^t = (v_{m,f}^t)_{m \in \mathcal{M}, f \in \mathcal{F}}$ with $v_{m,f}^t = \sum_{s \in \mathcal{S}^\Pi(m)} \lambda_{m,s,f}^t N_{s,f}$.

For a set of T Benders cuts (for some $T \geq 1$), we obtain an upper bound to the original problem CLCP by solving the surrogate IP

$$\begin{aligned} (\text{SUR-}T) \quad & \max_{\substack{(\mathbf{x}, \mathbf{z}) \in \mathcal{X} \\ \gamma \in \mathbb{R}_{\geq 0}}} \quad \gamma - \sum_{m \in \mathcal{M}} q_m z_m \\ \text{s. t.} \quad & \gamma \leq \Gamma^t + (\mathbf{v}^t)' \mathbf{x} \quad t = 1, \dots, T. \end{aligned}$$

Note that the auxiliary variable γ , together with the Benders cuts, approximates Slave linearly. This way, SUR- T avoids the non-linearity which creates the difficulty for solving Master. The optimal objective value to the surrogate problem is denoted by SUR T . The optimal solution consists of \mathbf{x}^{T+1} , \mathbf{z}^{T+1} and γ^{T+1} .

Generalized Benders decomposition is an iterative process. We start in step 0 with an initial feasible tuple of leasing and placement vectors $(\mathbf{x}^0, \mathbf{z}^0) \in \mathcal{X}$ and without any Benders cuts. At the start of step $T \geq 0$, we have the current leasing and placement $(\mathbf{x}^T, \mathbf{z}^T) \in \mathcal{X}$ and T Benders cuts. We solve Slave with input \mathbf{x}^T to obtain the optimal association vector \mathbf{y}^T . Clearly, since the triple \mathbf{x}^T , \mathbf{z}^T and \mathbf{y}^T are feasible to the original problem CLCP, its corresponding objective value provides a lower bound to the optimal value of CLCP. Additionally, we

compute the Lagrangian multipliers $\boldsymbol{\lambda}^T$ and the corresponding $(T+1)$ th Benders cut (5). With the Benders cut we obtain the surrogate MIP SUR- $(T+1)$. Its optimal solution is the feasible leasing and placement vectors $(\mathbf{x}^{T+1}, \mathbf{z}^{T+1}) \in \mathcal{X}$. The objective value SUR $^{T+1}$ is an upper bound to CLCP. This process iterates. At any step T , SUR T is the current upper bound (note that SUR $^{T+1} < \text{SUR}^T$ after every step T), while the current lower bound is provided by the best found solution $\max_{t \in \{0, \dots, T\}} h(\mathbf{y}^t) - \sum_{m \in \mathcal{M}} q_m z_m^t$. The process terminates in the globally optimal cache leasing and content placement vectors \mathbf{z}^* and \mathbf{x}^* when the upper and lower bounds coincide. Convergence is guaranteed from the proof of Theorem 2.4 in [18] and the fact that the domain \mathcal{X} of the Master is finite. In every step T , an instance of Master needs to be solved.

For general parameters Γ^t and \mathbf{v}^t , $T = 1, \dots, T$, it is NP-complete. This can be shown with a reduction from the classic SET COVER problem. State of the art MIP solvers such as CPLEX are, however, capable of solving SUR- T in reasonable runtime. We state without proof (due to space limitations) that the SUR- T problem has an infinite integrality gap (gap between linear relaxation and optimal discrete solution). Approximation algorithms for SUR- T are a topic for future research.

V. SPECIAL CASES

The savings function h introduced in Section III-D maps the user association vector \mathbf{y} onto the savings $h(\mathbf{y})$ (in €). In the following, we give examples of specific expressions for $h(\cdot)$ and explain what each example implies for the solution. We denote the cache-hit traffic volume associated to CBS m through vector \mathbf{y} by

$$v_m(\mathbf{y}) = \sum_{s \in \mathcal{S}^\Pi(m)} \sum_{f \in \mathcal{F}} y_{m,s,f}$$

A. Linear Savings (case i)

As a first case, the cache-hit user traffic of CBSs is linearly mapped onto monetary benefits for the CP, i.e.

$$h^L(\mathbf{y}) = c \sum_{m \in \mathcal{M}} v_m(\mathbf{y}), \quad (6)$$

where c is the savings (in €) per cache hit. Since the popularity vector is constant, the savings value is proportional to the hit ratio. The latter is simply calculated by $h^L(\mathbf{y}) / (c \sum_{s \in \mathcal{S}^\Pi} \sum_{f \in \mathcal{F}} N_{s,f})$.

With a linear savings function, the slave problem becomes easily tractable for any association policy Π , including CLOST and OPT-h: Given a CP vector \mathbf{x} , the MNO can freely distribute users among the covering CBSs $\mathcal{M}(s)$ which have f cached. The association to any CBS contributes equally to the savings. If no $m \in \mathcal{M}(s)$ caches f , then $y_{m,s,f} = 0$, hence no cache hit from region s for file f .

B. Separable Concave Savings (case ii)

As a second case, we introduce the sum of strictly concave functions, one per CBS, taking as argument the associated traffic volume. The strict concavity of the utility functions

implies diminishing returns for user traffic in every CBS. This way, the MNO has the incentive to associate users with underused CBSs while the overuse of CBSs is disincentivized. This choice for h can model physical resource limitations on each CBS that prohibit the good service of users when their volume becomes very high. As a result, a type of load balancing among CBSs is enforced. Observe that in this way, we do not impose a hard constraint on the user traffic but introduce a soft constraint in the objective function instead.

Formally, we define utility functions $U_m(\cdot)$ for every $m \in \mathcal{M}$ which are monotonously increasing, strictly concave and continuously increasing. The input of the utility functions is the cache-hit traffic volume at the CBS. The savings function is the sum of all utility functions. We obtain

$$h^c(\mathbf{y}) = \sum_{m \in \mathcal{M}} U_m(v_m(\mathbf{y})). \quad (7)$$

The load can be balanced among the CBSs by guaranteeing fairness with regards to associated volume. Some notions of fairness are max-min, α - and proportional fairness. Each of them is achieved by appropriate choice of the utility functions (see [10], [20]). E.g. for proportional fairness, utilities are chosen as (weighted) logarithms.

For any association policy Π , the slave problem with savings function as in (7) is a convex problem. Particularity the case $\Pi = \text{OPT-h}$ was solved in [13] using Lagrangian duality. The same method can be applied for CLOSEST as well.

C. Weighted Savings

In the previous section, there was no differentiation between wireless users in the objective function. Here, we introduce weights $w_{m,s,f} \geq 0$ that are specific to users from region s requesting content f associated to CBS m . This generalization allows to include costs and benefits from associating certain user groups to particular stations. The weighted traffic volume at CBS m is

$$v_m^{\mathbf{w}}(\mathbf{y}) = \sum_{s \in \mathcal{S}^\Pi(m)} \sum_{f \in \mathcal{F}} w_{m,s,f} y_{m,s,f},$$

where \mathbf{w} is the vector of weights. This volume can be used as argument of the weighted savings function (linear or concave)

$$h^{c,\mathbf{w}}(\mathbf{y}) = \sum_{m \in \mathcal{M}} U_m(v_m^{\mathbf{w}}(\mathbf{y})) \quad (8)$$

where c is the savings (in €) per weighted cache hit.

i) Prioritized Caching: If the weights $w_{m,s,f}$ are proportional to the file sizes b_f , files of larger size that create more burden to the backbone are favored to be cached. Such cases are also treated in [21].

ii) Network Throughput: The weights can represent the downlink throughput between a user and a CBS. For such a weights choice, the channel quality between m and s is a constant value $h_{m,s}$ that depends on a reference distance and the path loss exponent. The emitted power level of m is

denoted by p_m and the noise level by σ^2 . Then, the signal-to-interference-plus-noise ratio (SINR) of users in region s when associated to covering CBS m is

$$\text{SINR}_s(m) = p_m h_{m,s} \left(\sum_{\substack{\tilde{m} \text{ covers } s \\ \tilde{m} \neq m}} p_{\tilde{m}} h_{\tilde{m},s} + \sigma^2 \right)^{-1}, \quad (9)$$

where we assume that the interference from CBSs not covering s is negligible. For the downlink transmission from CBS m to region s , the throughput is equal to $w_{m,s,f} = B \log_2(1 + \text{SINR}_s(m))$ [in bits/sec] where B [Hz] is a chunk of bandwidth allocated to each served user. The total service bandwidth per CBS is equal to the product of B times the users routed to the CBS. Using these weights, CLCP takes into account that it is favorable for a CBS to serve users with good radio conditions in order to use its resources effectively. Such cases are also treated in [22].

We would like to emphasize that the arbitrarily many levels of channel quality on the coverage area of each station can be introduced by appropriately redefining the region set \mathcal{S} that partitions the network area. The modelling tradeoff is between the precision of communications aspects and the runtime of the optimization process.

VI. EXPERIMENTS AND NUMERICAL EVALUATION

A. Environment

We simulate cellular networks in an urban environment and calculate the optimal cache leasing and content placement for 4 cases: The savings function h is i) linear as in (6) or ii) the sum of utility functions as in (7). In the latter case all utility functions U_m are chosen as the natural logarithm to achieve proportional fairness for the user traffic at the CBSs. For each of the two cases, the MNO's user association policy is a) the MNO-CP cooperative policy OPT-h or b) the conventional policy CLOSEST. In each case, we simulated 100 random sets of CBS positions as a Poisson Point Process (PPP). This means that their total number in each run is a random Poisson realization, and their positions are uniformly distributed in the simulation window. The density of the PPP is $80 \frac{\text{CBS}}{\text{km}^2}$ for the cases i.a) and i.b) (linear savings) and $60 \frac{\text{CBS}}{\text{km}^2}$ for the cases ii.a) and ii.b) (log-savings). This implies an average minimal distance of 56m and 65m between the CBS positions, respectively. For the cases i.a) and i.b), the evaluation window has size $500 \times 500 \text{m}^2$, while the cases ii.a) and ii.b) were evaluated in a $300 \times 300 \text{m}^2$ window. The expected number of CBSs in the evaluated windows is 20 for case i) and 5.4 for case ii). In both cases, a larger area was simulated to avoid edge effects. The CBSs's coverage radius varies from 40m to 120m. The MNO price per unit size cache memory at all CBSs varies (0.01€-2.00€). The user population is distributed uniformly over the network with a density of 30 users per km^2 . The total simulated file catalog contains 100 objects. The content popularity follows the Zipf distribution with parameter 0.6 unless explicitly stated otherwise. The available cache size from the MNO is set to the catalog size so that only the pricing influences cache leasing decisions.

B. Implementation

All simulations have been performed using a native JAVA simulation environment. User association corresponding to the solution of the slave problem in Section IV is entirely done by optimization algorithms that we developed in the lab, outlined in Section V. The surrogate problem SUR- T in Section IV is solved using the state of the art mixed-integer problem solver IBM CPLEX 12.7.0 in combination with IBM ILOG CPLEX Optimization Studio. The experiments have been performed on a machine with a 2.40 GHz 16-core processor and 48 GB RAM.

C. Results for Linear Savings Function (case i)

At first, we present the simulation results for linear savings function. We emphasize case i.a) which performs OPT user association and compare it with case i.b) CLOSEST. Figure 2(a) illustrates how the hit ratio in case i.a) depends on the price per cache unit for different coverage radii. For all radii, the hit ratio decreases with increasing prices. With lowest price (0.01€/Unit), the CP leases in each CBS the entire memory available, so the hit ratio is 100%. As the price increases, the CP leases less units, and the hit ratio is reduced. This happens more quickly in networks with smaller coverage areas because there are less users covered by each CBS and also less coverage overlap area. When the price reaches a high level (2€), the cost from leasing cache memory exceeds the benefit from cache hits and the hit ratio drops to 0% for all coverage radii. The differences between the curves in Figure 2(a) diminish with higher radii where multi-coverage is already high enough.

The CP's leasing and placement decisions for networks with coverage radius 100m for the policies OPT-h (i.a) and CLOSEST (i.b) are shown in Figure 2(b). For each price, there are two columns: The lefthand-side column represents the i.a) case, the righthand-side column i.b). The height of each column is the average amount of cache units per CBS which are leased for the respective price. The subdivisions of each column represent the popularity of the files stored in the leased memory space: The bottom part are the ten most popular files, the second-to-bottom part are the files of popularity rank 11 to 20 and so on. For the lowest cache price (0.01€), all 100 available units are leased: For both assignment policies, the amount of leased cache memory decreases with increasing price. For the price interval 0.03-0.07€, OPT uses fewer cache units than CLOSEST. But, as the next Figure 2(c) shows, the hit probability by using OPT when the coverage radius is 100m and for the same price interval is 5-15% higher than the one achieved by CLOSEST even with less cached objects. Between 0.1 and 0.25€, OPT uses more cache memory than CLOSEST. However, the corresponding hit ratio is between 20-50% higher as well. For all prices, Figure 2(b) shows that the less popular files are represented more frequently with OPT than with CLOSEST, especially for prices ≥ 0.1 €. There is more diversity of visible content with the OPT placement.

Figure 2(c) directly compares OPT with CLOSEST. For all prices the hit ratio achieved by OPT (case i.a) is higher than

the one achieved by the CLOSEST policy (case i.b). In the price interval 0.06-0.50€, the relative hit ratio differences are higher when the coverage area of the CBSs is higher. For higher coverage radii (from 80m on) a hit ratio gap of 15%-50% can be seen. For higher prices, the CLOSEST hit ratio is close to 0, therefore the relative differences can become very high.

While Figure 2 shows the caching benefits for the CP, the costs it has to pay in return to the MNO (in case i.a) are depicted in Figure 3(a). The CP costs equal the MNO income. This amount can be calculated by multiplying the number of leased cache units with the price per unit. The maximum of the curve can be clearly identified for each radius. This is the operational point for the MNO when the latter aims for maximum income. The maxima are higher for larger coverage areas, while the difference in income decreases with increasing radius. Furthermore, the higher the coverage radius, the higher the cache leasing price at which the maximum is achieved.

The relation between hit ratio and MNO income can be seen in Figure 3(b): The x-axis displays the hit ratio achieved, the y-axis shows the income the MNO earns. Again, the income is higher in networks with larger coverage area. The maximum income for all simulated networks can be found for an achieved hit ratio between 80-90%. Conversely, if the CP decides to invest a certain sum, it can maximally achieve the rightmost of the two corresponding hit ratio values under the condition that the MNO chooses the pricing strategy most favorable to the CP.

The experimental results presented until now are based on a Zipf parameter of 0.6. However, a varying Zipf parameter influences the results: Figure 3(c) shows that the higher the price, the lower the hit ratio for any Zipf parameter in case i.a). This is due to the fact that lower price implies more leased units for the CP. For all prices (except the lowest one which achieves a hit ratio of near 100% throughout), the hit ratio increases with increasing Zipf parameter: With higher Zipf parameter, the population share requesting the most popular files becomes higher, thus caching popular files becomes more profitable and a higher hit ratio is achieved. Also, for the same price, the leased cache memory is more effective with higher Zipf parameters. The relative difference on hit ratio between cases i.a) (OPT) and i.b) (CLOSEST) depend on the Zipf parameter as well (not included in the paper). The lower the Zipf parameter, the more pronounced the differences in hit ratio between different cache prices since the benefits from overlapping coverage are bigger when content popularity is more even.

D. Results for Log Savings Function per CBS (case ii)

Here, we present the experimental results of case ii) in which the savings function is the sum of logarithms. Figure 4(a) shows the hit ratio for varying cache unit price both for the OPT (different coverage radii) and the CLOSEST policies. Due to the specific choice of the logarithmic savings function of case ii) the CLOSEST association gives identical cache leasing and content placement results for all coverage radii. For every coverage radius and every cache unit price,

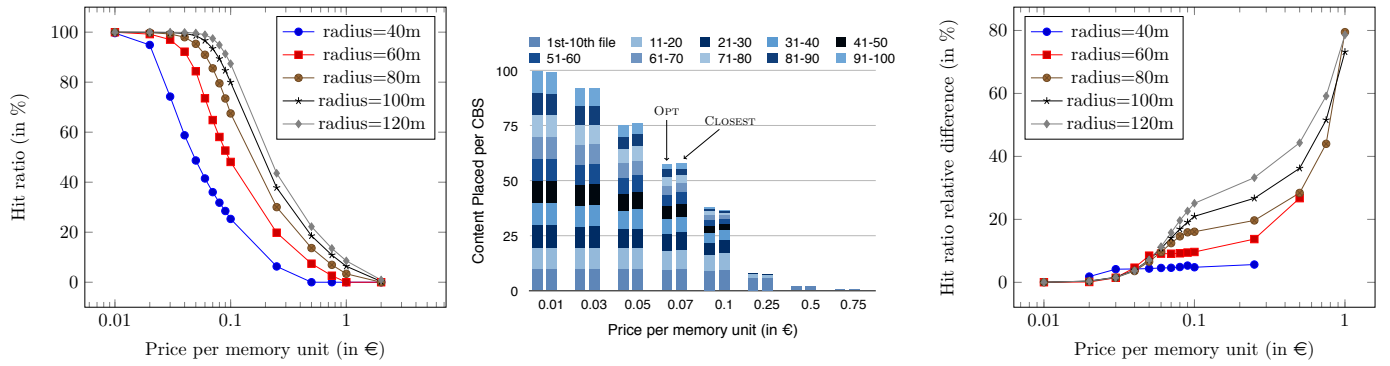


Fig. 2: Linear savings function. (a): Hit ratio in relation to price for different coverage radii for case i.a). (b): Cache lease and placement of popular files depending on cache unit price. Each left column represents case i.a), each right column is case i.b). (c): Relative difference between the hit ratio achieved by OPT-h (case i.a) to CLOSEST (case i.b) over cache unit price for different coverage radii.

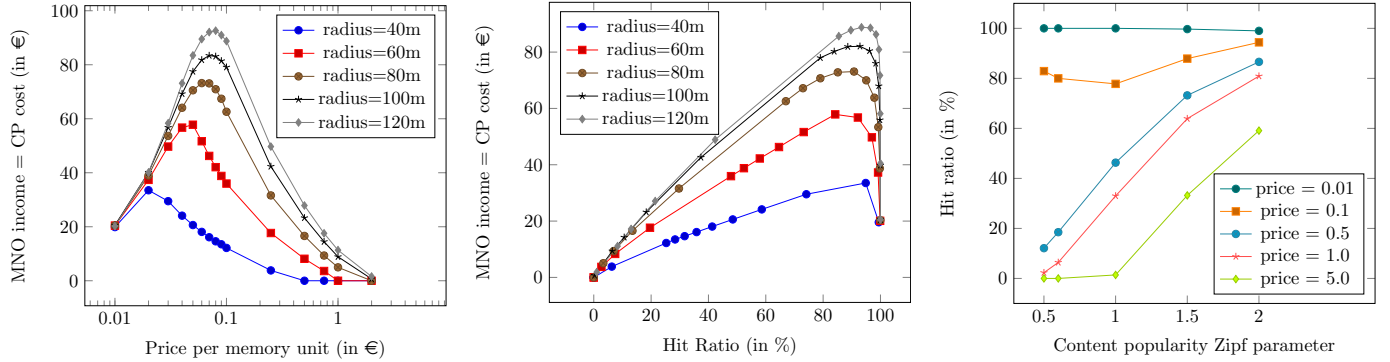


Fig. 3: Linear savings function. (a): Income of MNO in relation to price per cache unit for different coverage radii in case i.a). (b): MNO income/CP investment over hit ratio for different coverage radii in case i.a): Achieving the hit ratio on the x-axis results in the MNO income on y-axis. (c): Hit ratio in relation to Zipf parameter for different prices in case i.a).

the OPT policy achieves a higher hit ratio than the CLOSEST policy. Furthermore, the higher the coverage radius in case ii.a), the higher the hit ratio. The hit ratio improvement can reach over 15 percentage points using OPT. With increasing prices, caching becomes less profitable and the hit ratio decreases.

The advantage to the hit ratio of the OPT policy (ii.a) can be explained by the optimal content placement shown in Figure 4(b). Each pair of columns represents cache leasing and content placement for a certain unit price. Each left column represents the decisions according to OPT, each right column the decisions according to CLOSEST. The height is the average amount of cache units leased per CBS. The inner sections of the columns represent the content placement in all of the CBSs: The lowest section are the 10 most popular files, the second lowest the files ranked 11 to 20 and so on. It can be seen that particularly for low cache unit prices, the diversity of cached content is higher in case ii.a) than in case ii.b). For the lowest price (0.01€), CLOSEST provides only content from the more popular half of the catalog, while OPT-h places content from the tail of the catalog as well.

The main purpose of choosing the specific savings function in ii.a) is, however, the balancing of traffic load among the CBSs in order to avoid excess of resources by user

overflow which will lead to service dissatisfaction. Figure 4(c) shows that the additional load (from the increase in hit ratio using OPT, see Figure 4(a)) is distributed to the less loaded CBSs. The two upper (solid) lines in the graph represent the maximum load of a CBS in relation to the overall covered population per CBS both in the cases ii.a) and ii.b). The two lower (dashed) lines are the minimum loaded CBS. The maximum loaded CBSs in both ii.a) and ii.b) coincide as the figure shows. The minimum loaded CBS of ii.a) is higher than the ii.b), showing that excess users coming from the higher hit ratio are associated to the less loaded stations.

The three plots show that the OPT policy achieves an increase in hit ratio (good for both the CP and the MNO) while at the same time diversifying the cached content (good for the user) and avoiding an overload of CBSs (good for everybody).

VII. CONCLUSIONS

In this work, we propose a business model in which an MNO leases edge caches to a CP. The CP's objective is to maximize its savings through offloading of traffic from its data centers to the wireless caches while limiting the cache leasing costs. The optimality of the CP decisions depends on the MNO's user association policy as well as its pricing strategy.

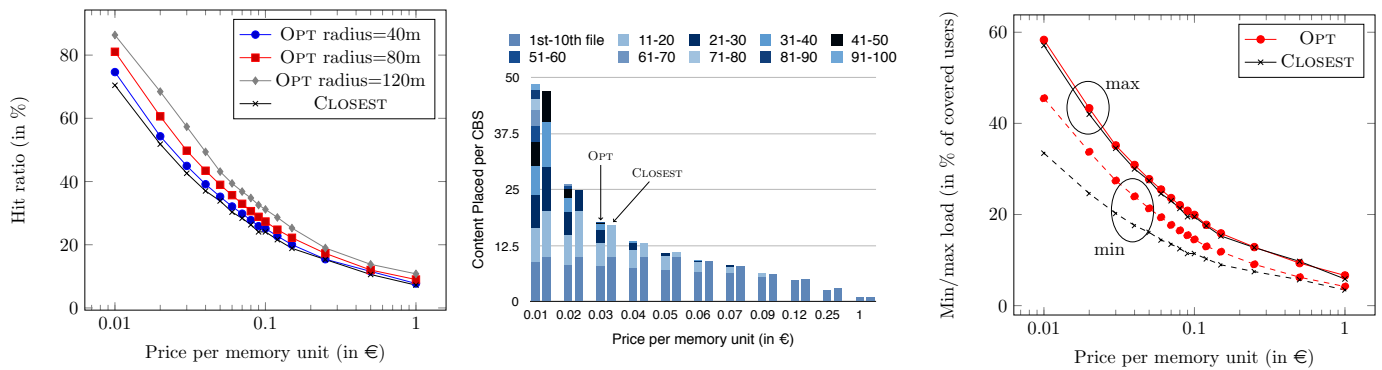


Fig. 4: Logarithmic savings function. (a): Hit ratio in relation to price for different coverage radii for cases ii.a) and ii.b). (b): Cache lease and placement of popular files depending on cache unit price. Each left column represents case i.a), each right column is case i.b). (c) Cache lease and placement of popular files depending on cache unit price. Each left column represents case i.a), each right column is case i.b).

The problem is modelled as a NLMIP taking the perspective of the CP. Network topology, association policy, pricing as well as CP savings function are general. Radio conditions and wireless node resource constraints can be included as soft constraints. An optimal solution to the general problem is found by applying Benders decomposition. The solution method converges to the global optimum and allows for each of the players to take separate actions iteratively.

Extensive experiments for random network topologies allow to compare the optimal CP decisions for different MNO association policies, cache prices, as well as CP savings functions. In all versions of the problem, we have identified a unique price that maximizes the MNO revenue. It depends on how much the CP valorizes traffic offloading achieved by the edge caches. This information is included in the CP's choice of the savings function. Another main conclusion is that MNO association policies that adhere to CP actions and exploit multi-coverage opportunities achieve higher offloading benefits for a given monetary investment. All these results suggest that the CP and MNO can jointly develop cooperative business models related to caching, that lead to considerable economic as well as operational benefits for both parties.

REFERENCES

- [1] Cisco visual networking index: Global mobile data traffic forecast update, 2015–2020 white paper. White Paper, 2 2016.
- [2] K. Shanmugam, N. Golrezaei, A.G. Dimakis, A.F. Molisch, and G. Caire. Femtocaching: Wireless content delivery through distributed caching helpers. *Information Theory, IEEE Transactions on*, 59(12):8402–8413, Dec 2013.
- [3] B. Błaszczyszyn and A. Giovanidis. Optimal geographic caching in cellular networks. In *Communications (ICC), 2015 IEEE International Conference on*, pages 3358–3363. IEEE, 2015.
- [4] E. Baştuğ, M. Bennis, and M. Debbah. Cache-enabled small cell networks: Modeling and tradeoffs. In *ISWCS 2014*, pages 649–653, Aug 2014.
- [5] K. Poularakis, G. Iosifidis, I. Pefkianakis, L. Tassiulas, and M. May. Mobile data offloading through caching in residential 802.11 wireless networks. *IEEE Transactions on Network and Service Management*, 13(1):71–84, 2016.
- [6] V. G. Douros, S. E. Elayoubi, E. Altman, and Y. Hayel. Caching games between content providers and internet service providers. *Performance Evaluation*, 2017.
- [7] A. Araldo, G. Dan, and D. Rossi. Stochastic dynamic cache partitioning for encrypted content delivery. In *Internet Teletraffic Congress (ITC) 2016*, 2016.
- [8] K. Poularakis, G. Iosifidis, and L. Tassiulas. Approximation algorithms for mobile data caching in small cell networks. *IEEE Transactions on Communications*, 62(10):3665–3677, 2014.
- [9] J. Roberts and N. Sbihi. Exploring the memory-bandwidth tradeoff in an information-centric network. In *Proceedings of the 2013 25th International Teletraffic Congress (ITC)*, pages 1–9, Sept 2013.
- [10] F. Kelly. Charging and rate control for elastic traffic. *Transactions on Emerging Telecommunications Technologies*, 8(1):33–37, 1997.
- [11] M. Dehghan, A. Seetharam, Bo Jiang, Ting He, Th. Salonidis, J. Kurose, D. Towsley, and R. Sitaraman. On the complexity of optimal routing and content caching in heterogeneous networks. In *INFOCOM*, 2015.
- [12] K.P. Naveen, L. Massoulié, E. Baccelli, A. Carneiro Viana, and D. Towsley. On the interaction between content caching and request assignment in cellular cache networks. In *ACM, AllThingsCellular '15*, pages 37–42. ACM, 2015.
- [13] J. Krolikowski, A. Giovanidis, and M. Di Renzo. Fair distributed user-traffic association in cache equipped cellular networks. In *WiOpt-CCDWN*, pages 1–6, 2017.
- [14] D. Bertsimas and R. Weismantel. *Optimization over integers*. Athena Scientific, 2005.
- [15] G. Paschos, E. Baştuğ, I. Land, G. Caire, and M. Debbah. Wireless caching: Technical misconceptions and business barriers. *IEEE Communications Magazine*, 54(8):16–22, 2016.
- [16] T. Bektaş, J. Cordeau, E. Erkut, and G. Laporte. Exact algorithms for the joint object placement and request routing problem in content distribution networks. *Computers & Operations Research*, 35(12):3860–3884, 2008.
- [17] A. Schrijver. *Theory of Linear and Integer Programming*. John Wiley & Sons, Inc., New York, NY, USA, 1986.
- [18] A. M. Geoffrion. Generalized Benders decomposition. *Journal of Optimization Theory and Applications*, 10:237–260, 1972.
- [19] A. Elwalid, D. Mitra, and Q. Wang. Cooperative data-optical internet-working: Distributed multi-layer optimization. In *INFOCOM*, 2006.
- [20] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Trans. Netw.*, 8(5):556–567, October 2000.
- [21] G. Neglia, D. Carra, and P. Michiardi. Cache policies for linear utility maximization. In *INFOCOM 2017, May 1-4, 2017*, pages 1–9, 2017.
- [22] J. Liu, B. Bai, J. Zhang, and K. B. Letaief. Cache placement in fog-rans: From centralized to distributed algorithms. *IEEE Transactions on Wireless Communications*, 16(11):7039–7051, Nov 2017.