

Joint Content Placement and Lightpath Routing and Spectrum Assignment in CDNs over Elastic Optical Network Scenarios

Jordi Perelló⁽¹⁾, Krzysztof Walkowiak⁽²⁾, Mirosław Klinkowski⁽³⁾, Salvatore Spadaro⁽¹⁾,
and Davide Careglio⁽¹⁾

- (1) Advanced Broadband Communications Center (CCABA), Universitat Politècnica de Catalunya (UPC),
Jordi Girona 1-3, 08034 Barcelona Spain, e-mail: perello@ac.upc.edu
- (2) Department of Systems and Computer Networks, Wrocław University of Technology, Poland
- (3) Department of Transmission and Optical Technologies, National Institute of Telecommunications
(NIT), Warsaw, Poland

Abstract: In this work, we address the problem of jointly deciding the placement of the contents delivered by a Content Distribution Network (CDN) among the available data centers, together with the allocation of the lightpaths required to serve the anycast demands initiated by the CDN network users, assuming an underlying high-capacity Elastic Optical Network (EON). We firstly present an Integer Linear Programming (ILP) formulation to optimally solve the targeted problem. This ILP formulation is of high complexity, though, and cannot be used to solve realistically sized problem instances. Hence, we also introduce a novel heuristic called CPRMSA-PD, which decomposes the problem into three sub-problems and applies greedy heuristics and simulated annealing meta-heuristic techniques to yield accurate solutions with practical execution times. We validate the performance of our CPRMSA-PD heuristic in medium-sized problem instances by comparing its results to the ones of the optimal ILP formulation. Next, we use it to give extensive insights into the effects of different key parameters identified in large CDN over EON backbone networks.

1. INTRODUCTION

The content-oriented network (CON) approach – named also information-centric networking, content-centric networking, content-aware network – has been recently proposed to facilitate the distribution of various contents (such as multimedia and data files) over the Internet and to provide a network infrastructure service, which is better suited to contemporary user needs and more resilient to disruptions and failures [1][2]. A CON solution that is currently used in the Internet is a content-delivery network (CDN), which is defined as a large distributed system of dedicated replica servers located in data centers (DCs), placed strategically across the Internet and providing replication of the content. The storage capacity of DCs is limited and it may be not possible to replicate the entire content in all DCs, thus some frequently accessed part of the content may be replicated in a larger number of DCs than a less popular one. In the context of CDNs, a new routing scheme called *anycast* routing appears. The *anycasting*, also defined as *one-to-one-of-many*, allows the customer to be connected to one of the multiple

DCs that provide requested content according to a certain criteria like cost, performance, resilience, etc. [3][4]. This approach assures high availability and improves the network resilience.

The latest Cisco report on Internet traffic forecast [5] predicts that CDNs will carry 62% of all Internet traffic by 2019, up from 39% in 2014. In terms of absolute values, the amount of CDN traffic in 2019 year will reach 84 EB per month, what is equivalent to permanent bit-rate of about 250 Tb/s. These extremely large numbers trigger the need to tackle the problem of content distribution related to CDNs directly in the backbone optical networks in order to provide good performance with high cost efficiency. In our paper, we focus on the recently proposed concept of Elastic Optical Networks (EONs). In more detail, the use of advanced, distance-adaptive transmission and modulation techniques, spectrum-selective switching technologies, and flexible frequency grids will allow next-generation optical transport networks to be spectrally efficient and, in terms of optical bandwidth provisioning, scalable and elastic when handling varying and ever increasing traffic demands when compared to traditional fixed-grid wavelength-switching optical networks (WSONs) [6]. Such flexible-grid EONs are considered as one of the technological pillars for building effective and cost-efficient content- and anycast-ready transport networks. Indeed, concentration of processing and storage capacity in relatively small number of DCs means that the volume of traffic on network links adjacent to these sites can become very large, thus network technologies supporting high capacity may be required. The high bandwidth scalability of EONs that utilize multi-carrier transmission within flexible frequency grids meets these requirements. Besides, the adaptation of transmission formats according to transmission path characteristics brings significant savings in terms of spectrum utilization [7]. Since the transmission over shorter routing paths can be performed with spectrally efficient modulation formats, the optical path (lightpath) connections carrying aggregated anycast traffic from nearer DCs to client nodes will occupy less spectrum resources than the lightpaths towards more distant DCs.

In this article, we focus on the problem of planning a CDN-oriented EON which, in principle, concerns the dimensioning and establishment of lightpath connections capable of delivering (by means of anycasting) the content replicated in DCs to the CDN users. This planning problem involves the problem of routing and spectrum allocation (RSA) – a basic optimization problem in EONs [8]. RSA consists in assigning to each lightpath a physical route in the network graph and a contiguous fraction of frequency spectrum subject to the constraint that the frequencies assigned to the connections do not overlap on the network links. RSA in EONs with anycast traffic demands was recently studied in the literature, for instance, in [4] and [7]. In all these works, an important simplification was assumed, which says that each DC hosts a replica of the entire content and a CDN user can be served by any DC. Moreover, it was assumed that the content provided in the CDN is of a single type. In fact, DCs have limited storage capacity that may lead to unequal replication and availability of content in DCs and, as a result, may influence the planning of lightpath connections. Moreover, the content could be of various types and with various popularity ratios, what should be addressed in the optimization process. In this work, we take this issue into account and we accept to have the content of various types unequally replicated in DCs.

The unequal replication of content involves another question, which is the assignment of a particular content to DCs. Consequently, a Content Placement (CP) problem arises. To address

such a CP, in this work we propose a model in which the entire content is divided into smaller Content Groups (CGs). Each CG is characterized by certain *popularity*, which represents an average relative rate at which CDN users try to access the content in that CG. Also, a CG can be replicated in any DC whenever there is available storage capacity. A reasonable and desirable outcome of our model is that more popular CGs are replicated in a larger number of DCs and, thus, placed closer to the CDN users, while less popular CGs are placed sparsely in the network. Eventually, all the content stored at a given DC and requested by users at a certain network node, even if belonging to different CGs, can be aggregated and delivered to the node users over a lightpath connection.

The inclusion of CP brings additional complexity into the lightpath planning problem since the volume of anycast traffic demands is not fixed anymore, as in the previous works, but depends on the placement of CGs and their availability in DCs. As a consequence, the previously proposed optimization models and algorithms for anycast RSA are not appropriate to the currently studied problem. Therefore, the main contributions of this article are threefold. First, we develop a novel optimization model for the joint CP and RSA problem in EONs, which is formulated as an Integer Linear Programming (ILP) model. Second, in order to solve the formulated problem, we propose a heuristic approach, where we decompose the problem into CP, traffic aggregation, and RSA sub-problems and solve each of these sub-problems using dedicated optimization methods. Finally, we perform a set of numerical experiments to evaluate the effectiveness of the heuristic and analyze the impact of different CDN parameters on spectrum requirements and transmission parameters of established lightpath connections.

The remainder of the paper is organized as follows. Section 2 presents related work in the literature. In section 3 we introduce the nomenclature to describe the network scenario and present a formal description of the optimization problem, which is formulated as an ILP problem in section 4. In Section 5 we describe a novel heuristic to solve the problem. In Section 6, we validate the performance of the heuristic against the ILP results, and use it to perform numerical experiments in large problem instances. Eventually, Section 7 includes concluding remarks.

2. RELATED WORK

The establishment of lightpaths in EONs involves the problem of RSA [6]. Additionally, if alternative optical signal modulation formats can be used, RSA is extended to the routing, modulation format selection, and spectrum assignment (RMSA) problem [9]. It was shown in several different ways that the RSA (thus also RMSA) optimization problem is *NP*-hard [8], [9], [10], [11], and several alternative ILP formulations [9], [12], [13], metaheuristics [14], [15], [16], and heuristics [8], [17], have been proposed to solve it.

An ILP formulation and a heuristic algorithm for RSA in EONs with anycast traffic demands have been proposed in [18] for an unprotected network and in [19] for a survivable network implementing a dedicated path protection scheme. Then, cost, energy and spectrum usage analysis of an EON supporting different modulation formats and serving anycast traffic demands has been performed in [7]. Moreover, the authors of [7] and [16] addressed the RMSA problem, which was formulated as an ILP problem with joint unicast and anycast flows, taking into account such optimization criteria as: CAPEX/OPEX network cost, power consumption, and spectrum usage. In both works, heuristic and metaheuristic algorithms were

proposed and evaluated. However it should be underlined that in all those works, the assumption was that each DC has an unlimited capacity and it contains the entire content.

Some works on CDNs have been focused on the phenomenon of content popularity. According to various analysis based on measurement of real Internet traffic, the content offered in CDNs follows different popularity by end users. This property allows designing the CDN in a more efficient way in order to improve overall performance. In particular, the authors of [20] and [21] present a traffic characterization study of the popular video sharing service, YouTube. In turn, in [22], it is demonstrated that web requests follow a Zipf-like distribution. The authors propose a model for web requests, where requests are independent and distributed according to a Zipf-like distribution. Reference [23] proposes a content caching scheme called WAVE. To reflect the content popularity, WAVE exponentially increases the number of chunks of a file with high popularity. According to presented results, WAVE achieves higher cache hit ratio and less frequent cache replacements than other on demand caching schemes that do not account for the content popularity property. Finally, the authors of [24] – using the observation that caching efficiency depends on the skewness of content popularity characteristics - propose a load balancing solution based on content popularity. Results presented in the paper show that the new approach outperforms mechanisms that are not aware of content popularity. For more general information on CDNs check [1], [2], [25], [26], [27] and references therein.

Moreover, there are some papers that address the problems of designing CDNs (e.g., [28], [29]) and content placement in CDNs (or equivalent problems), e.g., [30], [31], [32], [33], [34], [35]. However, these papers do not address routing of content requests or they assume a very simple routing model. Only few papers have considered a detailed routing of demands related to content provided in DCs (e.g., [36], [37]). However, the routing scenarios considered in these papers are based on classical multi-commodity flow modeling and do not address additional constraints of transparent optical networks, namely, the spectrum (wavelength) continuity constraint. Moreover, the authors of [36] and [37] design a disaster-resilient fixed-grid WSON in which a certain set of contents is given and the number of replicas of each content is limited, but not the storage capacity of DCs. Under similar assumptions, dynamic and disaster-aware CP in WSONs is addressed in [33]. To the best of our knowledge, there have been no papers that consider a joint optimization of content placement and routing and spectrum allocation in EONs addressing all constraints that occurs in transparent optical network and this paper is the first attempt to study this issue.

3. PROBLEM STATEMENT

In this section, we formally describe the problem that we aim to solve in this paper, namely, the joint Content Placement, lightpath Routing, Modulation format selection and Spectrum Assignment (CPRMSA) problem. We start by presenting the common nomenclature describing the network scenario (which will be used in subsequent sections) and then continue with the problem statement.

We assume the EON represented as an undirected graph $\mathcal{N} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} represents the set of network nodes (i.e., Bandwidth Variable Optical Cross Connects, BV-OXCs) and \mathcal{E} the set of unidirectional optical fiber links. Moreover, we assume that the spectrum of each fiber link is discretized as an ordered set of frequency slices $\mathcal{S} = \{s_1, s_2, \dots, s_{|S|}\}$. In this scenario, we pre-compute the set of candidate paths that the lightpaths can travel over the EON and we denote

it as \mathcal{P} . Regarding the Bandwidth Variable Transponders (BV-TXPs) to be equipped at network nodes, we assume without loss of generality that they can operate at a set of possible bit-rates \mathcal{B} . Moreover, they can employ a set of modulation formats \mathcal{M} at any bit-rate $b \in \mathcal{B}$.

To effectively account for the slice continuity and contiguity constraints that must be enforced for any lightpath throughout its route from source to destination, as well as its modulation format, we employ the candidate lightpath approach. Specifically, a candidate lightpath at a certain bit-rate $b \in \mathcal{B}$ employing modulation format $m \in \mathcal{M}$ is defined as a subset of adjacent slices along a path $p \in \mathcal{P}$ ensuring enough spectrum to allocate it. The number of such adjacent slices can easily be computed as $\lceil (b/ef_m + \Delta G)/W \rceil$ where ef_m is the efficiency of modulation format m (in bits/s/Hz), ΔG the required spectrum guard band between adjacent lightpaths (in GHz) and W the slice width (in GHz). Therefore, given $\mathcal{N}, \mathcal{P}, \mathcal{S}, \mathcal{B}, \mathcal{M}$ and ΔG , we can precompute the entire set of feasible candidate lightpaths (from any source to any destination) that we denote as \mathcal{L} , being \mathcal{L}_{uv} the subset of feasible candidate lightpaths from node u to node v , where $u, v \in \mathcal{V}$. Note that some candidate lightpaths may be unfeasible along a physical path due to insufficient transparent reach of the modulation format and, consequently, we do not include them into \mathcal{L} .

Regarding the CDN on top of the EON, we assume it composed of a set of DC replica servers \mathcal{R} , each one co-located with an EON node. We denote as n_r the capacity of $r \in \mathcal{R}$ (in generic capacity units, cu), and as a_r the EON node where the DC is located. Such replica servers can host the different content groups (CGs) into which the total content delivered by the CDN is divided. \mathcal{G} denotes the set of CGs offered by the CDN, and h_g and α_g identify, respectively, the size (in generic cu) and popularity of CG $g \in \mathcal{G}$. Note that, although the property of content popularity is in principle a dynamic feature, it is possible to estimate popularity of content groups in some period of time in order to obtain static values related to content popularity. This procedure can be repeated within some reasonable time scale in order to adapt the network to changing properties of the content. Finally, we denote as \mathcal{D} the set of offered anycast demands, and \mathcal{D}_g the subset of them requesting CG $g \in \mathcal{G}$. Moreover, the volume (in Gb/s) of anycast demand $d \in \mathcal{D}$ is represented as Γ_d , and the node initiating it as ∂_d . It is worth mentioning here, that in this work we will only account for the downstream traffic associated to each anycast demand in \mathcal{D} . Being totally aware that anycast demands are bidirectional in reality, their upstream traffic volume (from the client to the DC) is much lower and almost negligible against their downstream one (from the DC to the client). Therefore, we assume here that the low bit-rate of upstream demands could be served over the conventional TCP/IP Internet or any separated packet-switched network, while large downstream demands are carried over an EON.

This being said, the CPRMSA problem can be formally described as:

Find: the placement of CGs in the DC replica servers and candidate lightpaths allocated over the EON, subject to the following *constraints*:

- 1) *Content group placement:* all CGs must be placed in one DC replica server at least (e.g., more popular CGs can be replicated in several DC replica servers).
- 2) *Replica server capacity:* the sum of the sizes of the CGs hosted by a DC replica server cannot exceed its capacity.

- 3) *Anycast demand delivery*: all requested anycast demands must be successfully delivered (no blocking is assumed in this work). An anycast demand can be delivered from any DC replica server hosting the requested CG.
- 4) *Lightpath capacity*: the sum of the capacities of all lightpaths established from a DC replica server location to a certain network node must be enough to carry all anycast demands initiated from that network node and delivered by that replica server.
- 5) *Spectrum slice clashing*: a spectrum slice of a certain link can support one candidate lightpath at most.

with the *objective* to minimize the number of spectrum slices used at any link in the network, a typical optimization objective in existing works on the topic (e.g., in [8] and [9]).

To illustrate the optimization problem addressed in this paper, we consider a simple example shown in Figure 1. The network includes five nodes, where two of them host a DC (nodes D and E), while the three remaining nodes A, B, and C are client nodes (where CDN users are connected). The content requested by the clients is divided into four content groups CG1, CG2, CG3 and CG4. The figure shows a configuration of content placement. In particular, due to limited storage capacity, DC located at node D stores only content related to CG1, CG2, CG3 and DC located at node E stores only content related to CG2, CG3, CG4. For each client node, the requested bit rate related to a particular CG is shown. For instance, in the case of node A, the volume of demand related to CG1 is 70 Gb/s, the volume of demand related to CG2 is 50 Gb/s, etc. Note that the demands are anycast, if the requested CG is hosted in more than 1 DC node. In our example, this condition is fulfilled in the context of demands connected to CG2 and CG3. For each demand, after the bit rate value, the selected DC node is reported in parentheses, e.g, for node A CG1 is served by DC located at node D, CG2 is transmitted from node E. This assignment yields the aggregated volume of bitrate to be delivered from DCs to each client nodes, i.e., the bitrate from node E to node B is 130 Gb/s. To provision the aggregate traffic between each node pair, a dedicated lightpath is established. Figure 1 shows all created lightpaths in colors related to particular client nodes. The lightpaths are provisioned using the same range of spectrum slices throughout their route, what is presented in the figure for each link.

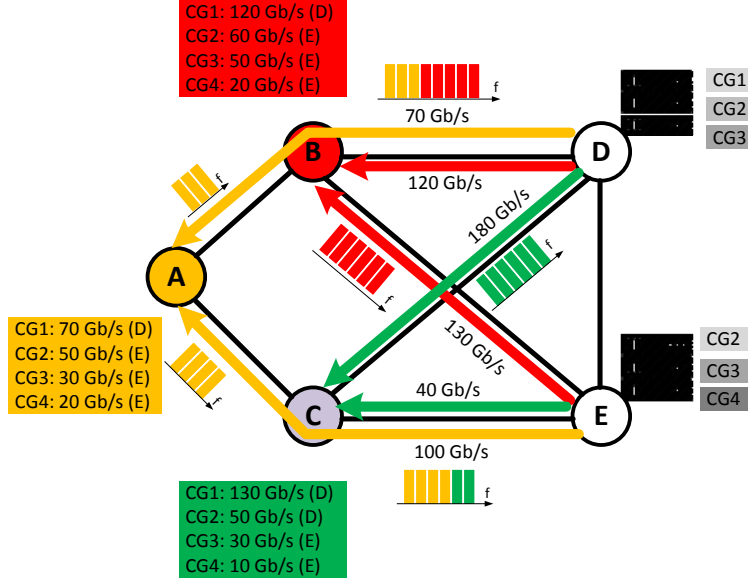


Figure 1. Example illustrating the problem of planning a CDN-oriented EON.

4. OPTIMAL CPRMSA-ILP FORMULATION

In this section we present a novel ILP formulation for the CPRMSA problem, hereafter referred as CPRMSA-ILP. To the best of our knowledge, no ILP formulation for a problem like this one has been proposed to date in any work in the literature.

The following decision variables have been defined:

z_{rg} : Binary; 1 if DC replica server $r \in \mathcal{R}$ hosts popularity group $g \in \mathcal{G}$; 0 otherwise.

q_{dr} : Binary; 1 if demand $d \in \mathcal{D}$ is served in DC replica server $r \in \mathcal{R}$; 0 otherwise.

f_{uv} : Real; aggregated flow from node u to node v , $u, v \in \mathcal{V}$.

x_l : Binary; 1 if candidate lightpath $l \in \mathcal{L}$ is established; 0 otherwise.

y_{es} : Binary; 1 if slice $s \in \mathcal{S}$ in fiber link $e \in \mathcal{E}$ is used; 0 otherwise.

y_s : Binary; 1 if slice $s \in \mathcal{S}$ is used in any fiber link of the network; 0 otherwise.

And the ILP formulation reads as:

$$\text{minimize } F = \sum_{s \in \mathcal{S}} y_s \quad (1)$$

subject to:

$$\sum_{r \in \mathcal{R}} z_{rg} \geq 1, \quad \forall g \in \mathcal{G} \quad (2)$$

$$\sum_{g \in \mathcal{G}} h_g z_{rg} \leq n_r, \quad \forall r \in \mathcal{R} \quad (3)$$

$$q_{dr} \leq z_{rg}, \quad \forall r \in \mathcal{R}, g \in \mathcal{G}, d \in \mathcal{D}_g \quad (4)$$

$$\sum_{r \in \mathcal{R}} q_{dr} = 1, \quad \forall d \in \mathcal{D} \quad (5)$$

$$f_{uv} = \sum_{r \in \mathcal{R}: a_r = u} \sum_{d \in \mathcal{D}: \partial_d = v} \Gamma_d q_{dr}, \quad \forall u, v \in \mathcal{V} \quad (6)$$

$$\sum_{l \in \mathcal{L}_{uv}} c_l x_l \geq f_{uv}, \quad \forall u, v \in \mathcal{V} \quad (7)$$

$$\sum_{l \in \mathcal{L}: e \in l \wedge s \in s_l} x_l = y_{es}, \quad \forall e \in \mathcal{E}, s \in \mathcal{S} \quad (8)$$

$$\sum_{e \in \mathcal{E}} y_{es} \leq |\mathcal{E}| y_s, \quad \forall s \in \mathcal{S} \quad (9)$$

Objective function (1) contemplates our optimization goal, that is, the minimization of the slices used in any network link. Constraint (2) enforces that every CG is placed at least in one DC replica server. Constraint (3) is the DC replica server capacity constraint, ensuring that the sum of the sizes of all hosted CGs does not exceed its capacity. Constraint (4) accounts for the DC replica servers that can serve a particular anycast demand, namely, those hosting the CG that the anycast demand requests. Constraint (5) enforces that each anycast demand is served by exactly 1 replica server. Constraint (6) builds the total aggregated flow from any node u to any node v , $u, v \in \mathcal{V}$, accounting for the DC replica servers from which anycast demands are delivered. Constraint (7) is the lightpath capacity constraint, ensuring that the sum of the capacities of the selected candidate lightpaths is enough to carry the total aggregated flow from a source to a destination node. Constraint (8) is the slice clashing constraint, accounting as well for the slices used in every network link. Finally, constraint (9) keeps the slices used along the entire network. From this description, we can easily compute the total number of decision variables of CPRMSA-ILP as $|\mathcal{R}||\mathcal{G}| + |\mathcal{D}||\mathcal{R}| + |\mathcal{V}|^2 + |\mathcal{L}| + |\mathcal{E}||\mathcal{S}| + |\mathcal{S}|$. Moreover, the total number of constraints is $|\mathcal{G}| + |\mathcal{R}| + |\mathcal{R}||\mathcal{D}| + |\mathcal{D}| + 2|\mathcal{V}|^2 + |\mathcal{E}||\mathcal{S}| + |\mathcal{S}|$.

Note that the CPRMSA problem is NP -complete in general. Indeed, consider a special case of CPRMSA, in which: (a) $|\mathcal{D}| = |\mathcal{G}| = |\mathcal{R}|$ and $|\mathcal{D}_g| = 1, \forall g \in \mathcal{G}$, i.e., the number of demands equals the number of CGs and the number of replica servers, and each CG is requested by only one demand, and (b) for each $g \in \mathcal{G}$, there exists one and only one $r \in \mathcal{R}$ that satisfies $h_g = n_r$; in other words, a feasible content placement solution requires CG g to be located in replica server r . Consequently, for each demand $d \in \mathcal{D}$ we have to establish a lightpath between the origin of d and determined replica server r hosting the CG requested by d . This special case of CPRMSA is equivalent with the RSA problem, which is NP -complete [8], [9], [10], [11].

In fact, CPRMSA is substantially more complex than RSA since it targets the joint optimization of the placement of the CGs in the DC replica servers, as well as the allocation of the candidate lightpaths carrying the anycast demands. Thus, given its high complexity, it will not be solvable for large problem instances, but will allow us assessing the performance of the CPRMSA-PD heuristic in moderately sized problem instances.

5. CPRMSA-PD HEURISTIC

The proposed CPRMSA-PD heuristic decomposes the CPRMSA problem into 3 sub-problems solved separately: 1) content placement; 2) demand aggregation; 3) candidate lightpath selection. Each of these sub-problems entails significantly lower complexity than the complete CPRMSA problem. Thus, solving them separately allows to reduce drastically the execution times, while achieving approximate results very close to the optimal ones (as will be shown later on in subsection 6.2). The general pseudo-code of the heuristic is presented in Figure 2.

Input: $\mathcal{N}, \mathcal{R}, \mathcal{G}, \mathcal{D}$, numOfGlobalIter, maxIter, $\beta, \gamma, T_{coef}, \tau$

Output: BestSol

1: BestSol = \emptyset


```

2:  $z_{rg} \leftarrow \text{ContentPlacement}(\mathcal{N}, \mathcal{R}, \mathcal{G}, \mathcal{D}, \beta)$ 
3: for glter = 1 to numOfGlobalter do
4:    $\Delta \leftarrow \text{DemandAggregation}(\mathcal{N}, \mathcal{R}, \mathcal{G}, z_{rg}, \mathcal{D}, \gamma)$ 
5:    $Sol \leftarrow \text{CandidateLightpathSelection}(\mathcal{N}, \Delta, \text{maxIter}, T_{coef}, \tau)$ 
6:   if  $F_{Sol} < F_{BestSol}$  then
7:      $BestSol = Sol$ 
8: return  $BestSol$ ;

```

Figure 2. General CPRMSA-PD pseudo-code.

Looking at the pseudo-code, CPRMSA-PD firstly solves the content placement sub-problem (line 2), storing the result in the output binary variables z_{rg} (same role as in the ILP formulation). Then, an iterative procedure (lines 3-7) is initiated, where the offered anycast demands are aggregated into unicast demands from DC locations to nodes where anycast demands are initiated (line 4). The resulting list of unicast aggregated demands, denoted as Δ , is fed as input to the candidate lightpath selection (line 5), which decides on the best candidate lightpaths to serve the entire list of unicast aggregated demands (Sol). If the number of slices used in the network (F_{Sol}) is lower than in the best solution encountered until the moment ($F_{BestSol}$), Sol is stored as $BestSol$. At the end of the $numOfGlobalter$ iterations, $BestSol$ is delivered as the solution of the heuristic (line 8).

In what follows, the procedures employed to solve each of the identified sub-problems are further elaborated.

5.1. Content placement sub-problem

This sub-problem is solved by means of a simple ILP formulation that aims at placing the different CGs at the available DC replica servers, while minimizing the cost of delivering the anycast demand flows to the clients and the maximum amount of flow carried by any EON fiber link (with load balancing purposes). Specifically, C_{rv} denotes the cost of delivering one unit of anycast demand flow from DC replica server $r \in \mathcal{R}$ to node $v \in \mathcal{V}$, which is set in this work to the number of hops of the shortest physical path from the node where replica server r is located (i.e., a_r) to v , identified as p_{rv} . For this ILP formulation, most of the nomenclature and decision variables presented in previous sections apply. Specifically, decision variable f_{rv} (real-valued) accounts for the total amount of demand flow from a_r to node $v \in \mathcal{V}$. Moreover, decision variable M (real-valued) keeps the maximum amount of flow carried at any EON link. Then, the ILP formulation is:

$$\text{minimize } \phi = \beta[\sum_{r \in \mathcal{R}} \sum_{v \in \mathcal{V}} C_{rv} f_{rv}] + (1 - \beta)M \quad (10)$$

subject to:

$$\sum_{r \in \mathcal{R}} z_{rg} \geq 1, \quad \forall g \in \mathcal{G} \quad (11)$$

$$\sum_{g \in \mathcal{G}} h_g z_{rg} \leq n_r, \quad \forall r \in \mathcal{R} \quad (12)$$

$$q_{dr} \leq z_{rg}, \quad \forall r \in \mathcal{R}, g \in \mathcal{G}, d \in \mathcal{D}_g \quad (13)$$

$$\sum_{r \in \mathcal{R}} q_{dr} = 1, \quad \forall d \in \mathcal{D} \quad (14)$$

$$f_{rv} = \sum_{d \in \mathcal{D}: \partial_d = v} \Gamma_d q_{dr}, \quad \forall r \in \mathcal{R}, v \in \mathcal{V} \quad (15)$$

$$\sum_{r \in \mathcal{R}} \sum_{v \in \mathcal{V}: e \in p_{rv}} f_{rv} \leq M, \quad \forall e \in \mathcal{E} \quad (16)$$

Objective function (10) has the twofold goal mentioned above, where the emphasis put into one or the other optimization objective is provided by parameter β , $0 \leq \beta \leq 1$. Constraints (11)-(15) have the same goals as constraints (2)-(6) in the CPRMSA-ILP. Finally, constraint (16) allows keeping the maximum flow carried at any network link, assuming that flows are carried over the shortest path from the DC locations and the nodes initiating the demands. As mentioned before, the values assigned to the binary decision variables z_{rg} are the output of this sub-problem.

5.2. Demand aggregation sub-problem

At first sight, it seems that delivering an anycast demand from the closest replica server to the client should minimize the spectrum slices used in the network. While possibly true when the sum of slices occupied over all network links is the minimization target, this may not hold in our case. Recall that we are aiming to minimize the number of slices used at any link in the network, which is very influenced by the excessive occupation of certain links (typically those adjacent to highly demanded DCs).

Figure 3 describes the proposed demand aggregation procedure. This procedure starts (line 1) with an empty list of aggregated demands (Δ). Next, it tries to aggregate all offered anycast demands, one per iteration (lines 2-18). For each anycast demand, the list of DCs hosting the desired CG is searched (line 3). Being only hosted in one DC, this DC is unavoidably selected as the one serving the demand (lines 4-5). Conversely, to better explore the solution space when the desired CG is hosted in > 1 DC, we employ parameter γ as the probability that an anycast demand is not served from the closest DC hosting the desired CG, but the next one in terms of hops of the shortest path to it (lines 6-11). As a result, Δ can differ through the global iterations depicted in the CPRMSA-PD heuristic pseudo-code.

Having decided the DC replica server from which an anycast demand is delivered, this one turns to be a unicast demand. In order to effectively fill the capacity of the lightpaths that will be allocated in the network, the proposed procedure tries to aggregate all anycast demands from the same DC to the same network node into aggregated demands of capacity lower or equal to the maximum capacity of a lightpath B_{max} (lines 12-14). If a demand cannot be added into the existing aggregated demands, a number of new aggregated demands from the same DC to the same node are created, providing enough additional capacity to carry the anycast demand (line 16). The capacity of the newly created aggregated demands is filled with the volume of the current anycast demand and stored in Δ (lines 17-18).

Input: $\mathcal{N}, \mathcal{R}, \mathcal{G}, z_{rg}, \mathcal{D}, \gamma$

Output: Δ

```

1:  $\Delta = \emptyset$ 
2: for each  $d$  in  $\mathcal{D}$  do
3:    $\Psi_d \leftarrow$  DC replica servers hosting CG requested by  $d$ 
4:   if size of  $\Psi_d = 1$  then
5:      $r \leftarrow \Psi_d[0]$ 
6:   else

```

```

7:      Sort every  $r \in \Psi_d$  according to the hops of the shortest path between  $a_r$  and
      node  $v \in \mathcal{V}$  where the client of the demand is attached, in increasing order
8:      if  $p < \gamma$  do
9:           $r \leftarrow \Psi_d[1]$ 
10:     else
11:          $r \leftarrow \Psi_d[0]$ 
12:     Find if the sum of available capacities of the aggregated demands from  $a_r$  to  $v \geq \Gamma_d$ 
13:     if true then
14:         Subtract  $\Gamma_d$  from the capacities of available aggregated demands
15:     else
16:         Create as many additional aggregated demands from  $a_r$  to  $v$  as necessary to
         carry  $\Gamma_d$ 
17:         Subtract  $\Gamma_d$  from the capacities of the available aggregated demands
18:         Add the newly created aggregated demands in  $\Delta$ 
19:     return  $\Delta$ ;

```

Figure 3. Demand aggregation procedure.

5.3. Candidate lightpath selection sub-problem

Finally, this sub-problem aims at selecting the best candidate lightpaths in \mathcal{L} to serve all aggregated demands in Δ , so that objective function F is minimized. As discussed before, the selection of a candidate lightpath implicitly selects its bit rate, modulation format and route and spectrum assignment. To this goal, a Simulated Annealing (SA)-based meta-heuristic is proposed, which guides a simpler candidate lightpath selection greedy heuristic. SA is a well-known probabilistic meta-heuristic method inspired from the annealing processes in metallurgy. In SA, a temperature parameter is initialized and cooled-down as the algorithm evolves. The higher this temperature, the more probable is to accept non-improving solutions, which allows SA escaping from local optima. This happens more frequently at the early stages of the algorithm, while only better solutions are generally accepted at the end (temperature values become low). The interested reader can find additional information on the SA meta-heuristic method in [38].

Figure 4 details the complete candidate lightpath selection procedure pseudo-code. At the beginning, aggregated demands in Δ are sorted according to their minimum slices required in descending order. For this, the number of slices required by the most efficient candidate lightpath among all candidate lightpaths available to serve the aggregated demand (L_δ) is taken into consideration (lines 1-3). Next, the simple candidate lightpath selection heuristic (detailed later on in Figure 5) is run to obtain an initial solution of the sub-problem (line 4), i.e., the best solution found so far. As starting temperature value for the SA (line 5), we set $T = F_{iniSol} * T_{coef}$, that is, the number of slices used in the initial solution of the greedy heuristic multiplied by the temperature coefficient (T_{coef}). To effectively explore the solution space, SA swaps a randomly chosen pair of aggregated demands in Δ per iteration changing their order (lines 8-10), and the greedy heuristic is run again processing the new reordered Δ (line 11). If the new solution uses a lower number of slices, the new order is kept for subsequent iterations. Even if a worse solution is found, SA decides to keep the new order with probability equal to $e^{-\Omega/T}$, i.e., the Boltzmann function, which allows SA to escape from local

optima. Otherwise, the previous swap operation is undone. All these operations are performed between lines 12-17. The current temperature is decreased at the cooling rate τ per iteration (line 18). Finally, the best found solution is returned as the output of the procedure (line 20).

Input: $\mathcal{N}, \Delta, \text{maxIter}, T_{coef}, \tau$

Output: BestSol

```

1: Pre-compute the set of candidate lightpaths  $\mathcal{L}_\delta$  for every aggregated demand
2: Find the minimum number of required slices by every aggregated demand
3: Sort  $\Delta$  according to the minimum number of required slices, in descending order

4: BestSol  $\leftarrow$  CandidateLightpathSelectionGreedyHeuristic ( $\mathcal{N}, \Delta$ )
5:  $T = F_{BestSol} * T_{coef}$ 

6: iter = 0
7: while iter < maxIter and  $T > 0$  do
8:    $\delta_1 =$  choose random aggregated demand from  $\Delta$ 
9:    $\delta_2 =$  choose random aggregated demand from  $\Delta$ , so that  $\delta_1 \neq \delta_2$ 
10:  swap( $\Delta, \delta_1, \delta_2$ );
11:  Sol  $\leftarrow$  CandidateLightpathSelectionGreedyHeuristic ( $\mathcal{N}, \Delta$ )
12:   $\Omega = F_{Sol} - F_{BestSol}$ 
13:  if  $\Omega < 0$  then
14:    BestSol = Sol;
15:  else
16:    if  $e^{-\Omega/T} \leq \text{rnd} < 1$  then //rnd = random value [0,1)
17:      swap( $\Delta, \delta_1, \delta_2$ ); //Restore original order
18:     $T = T * \tau$ 
19:    iter++;

20: return BestSol;

```

Figure 4. Candidate lightpath selection procedure.

To complete the procedure above, Figure 5 shows the pseudo-code of the candidate lightpath selection greedy heuristic used. This heuristic runs an iterative process, where at each iteration, the maximum reservable slice (*maxSlice*) in any link is set to its value in the previous iteration (0 at the beginning) plus the minimum required number slices of the first pending demand in Δ , i.e., the one requiring more slices (line 3). Next, for each pending aggregated demand $\delta \in \Delta$, the most efficient available candidate lightpath in L_δ with enough capacity is sought on a first-fit basis, never exceeding *maxSlice* (line 5). If an available candidate lightpath is found, the associated resources are reserved and the aggregated demand is considered as served (lines 6-8). The iterative process stops when all aggregated demands in Δ are successfully served (no blocking is considered in this work).

Input: $\mathcal{N}, \Delta, \mathcal{L}$

Output: Sol

```

1: maxSlice = 0
2: while any pending aggregated demand in  $\Delta$  do
3:     maxSlice += required slices of the first pending demand in  $\Delta$ 
4:     for each pending demand  $\delta$  in  $\Delta$  do
5:          $l$  = First available candidate lightpath in  $L_\delta$  to serve  $\delta$  (using slices 1..maxSlice)
6:         if  $l$  found then
7:             Reserve physical resources supporting candidate lightpath  $l$ 
8:             Mark  $\delta$  as served
9: return Sol

```

Figure 5. Candidate lightpath selection greedy heuristic.

5.4. CPRMSA-PD complexity analysis

Being presented the strategies to address the sub-problems into which CPRMSA-PD decomposes the complete CPRMSA problem, this subsection focuses on its computational complexity. From the general CPRMSA-PD pseudo-code (Figure 2), it requires a single execution of the proposed CP ILP, plus numOfGlobalIter executions of both demand aggregation and candidate lightpath selection heuristics.

Firstly, the total number of decision variables of the ILP formulation for the CP sub-problem becomes $|\mathcal{R}||\mathcal{G}| + |\mathcal{D}||\mathcal{R}| + |\mathcal{R}||\mathcal{V}| + 1$ and the total number of constraints is $|\mathcal{G}| + |\mathcal{R}| + |\mathcal{R}||\mathcal{D}| + |\mathcal{D}| + |\mathcal{R}||\mathcal{V}| + |\mathcal{E}|$. Such numbers are much lower than those in the complete CPRMSA-ILP (see Section 4), as neither the candidate lightpath selection nor the aggregation of anycast demands into them are taken here into account.

Focusing on the complexity of the procedure for the demand aggregation sub-problem, it is $\approx O(|\mathcal{D}||\mathcal{R}| \log|\mathcal{R}|)$ provided that an efficient ($n \log n$ complexity) sorting algorithm is used (in Figure 3, line 7). Next, regarding the candidate lightpath selection sub-problem, we can find that the complexity of the simpler greedy heuristic in Figure 5 is bounded by $O(|\Delta| \cdot [1 + |\Delta|] \cdot |L_\delta|/2) \approx O(|\Delta|^2 |L_\delta|)$ assuming the (improbable) worst-case when only a single aggregated demand can be allocated per iteration. As a result, the computational complexity of the overall SA-based candidate lightpath selection procedure (Figure 4) is bounded by $\approx O(\maxIter \cdot |\Delta|^2 |L_\delta|)$. Note that the pre-computation procedures in Figure 4 (lines 1-3) have not been considered in this complexity analysis, since they could be previously performed and read (e.g., from text files) upon starting CPRMSA-PD (they are shown in Figure 4, though, for better understanding of the CPRMSA-PD operation). All these values allow us to state the total complexity of CPRMSA-PD as that of the CP ILP plus $\approx O(\text{numOfGlobalIter} \cdot \maxIter \cdot |\Delta|^2 |L_\delta|)$.

6. RESULTS AND DISCUSSION

This section is aimed to assess the performance of the CPRMSA-PD heuristic for solving the CPRMSA problem against the optimal CPRMSA-ILP. To solve the latter, the commercial CPLEX v12.5 optimization software has been used. All executions are run in a commercial 8-core Intel i7 PC at 3.4 GHz with 16 GB RAM. Once CPRMSA-PD performance is validated, we use it to numerically quantify the effects of modifying different key parameters in the targeted

scenario, such as the number of DCs and capacity of the replica servers therein, number of CGs into which the whole content is split and CG popularity distribution.

6.1. Scenario details and assumptions

In order to generalize our conclusions as much as possible, three different network topologies have been evaluated (Figure 6): a medium-sized INT9 network (9 nodes and 26 unidirectional links) and two large reference backbone networks, namely, the US26 network (26 nodes and 84 unidirectional links) and Euro28 network (28 nodes and 82 unidirectional links). In all these scenarios, the entire C-Band 4 THz spectrum discretized in 12.5 GHz frequency slices has been made available, resulting into 320 slices. Concerning the BV-TXP equipped at EON nodes, we assume that they can operate either at 100, 200, 300 or 400 Gb/s. Moreover, they can employ any of the following modulation formats: Polarization Mode (PM)-BPSK, PM-QPSK, PM-8-QAM or PM-16-QAM, with spectral efficiencies of 2, 4, 6 and 8 bits/s/Hz, respectively (i.e., the PM capability doubles their initial spectral efficiency). As in [9], we assume their transparent reach to be 3000, 1500, 750 and 375km, respectively, independently of the actual bit rate.

To generate the candidate lightpaths over the EON, for each pair of source-destination nodes we compute the first K physical shortest paths ($K = 3$ in this work). For each of these shortest paths and available BV-TXP bit-rate, we then select the most efficient modulation format that can cover the entire path length. If no modulation format among the available ones can cover it transparently (i.e., due to insufficient transparent reach), we choose the modulation format that minimizes the number of required regeneration stages, assuming for simplicity that no spectrum conversion can be performed during regeneration. Next, we calculate the required number of adjacent frequency slices is computed using the expression previously presented in section 3, with $\Delta G = 10$ GHz. Finally, the candidate lightpaths at that bit-rate, using that modulation format and requiring that number of adjacent frequency slices are computed along that path, always ensuring the spectrum continuity constraint (selected adjacent slices cannot change along the links composing the path).

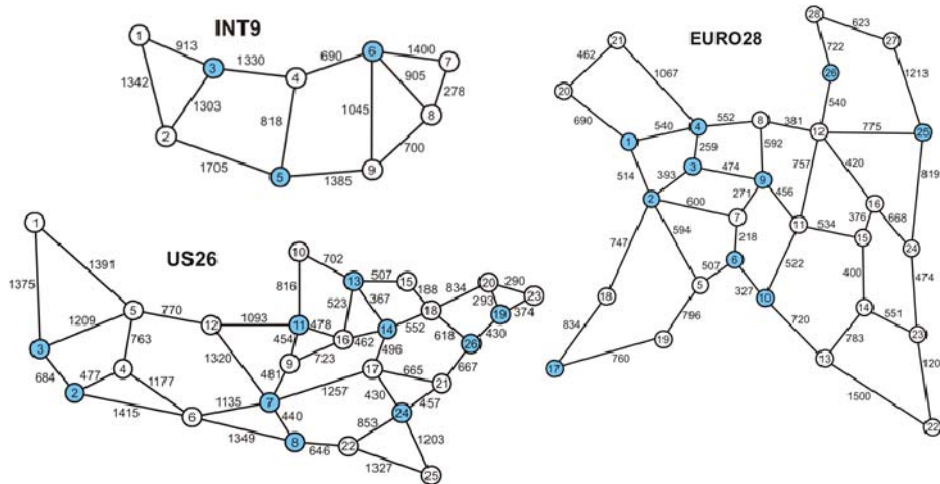


Figure 6. Network topologies used, where DC locations have been highlighted (in blue). Link distances are shown in km.

We assume 2 scenarios for each network topology, with 2 and 3 DCs in the INT9 network, and 5 and 10 DCs in both US26 and Euro28 backbone networks. Table 1 provides further information on the location of these DCs, which have also been highlighted in Figure 6.

Network topology	$ \mathcal{R} $	Node locations
INT9	2	3,6
	3	3,5,6
US26	5	3,8,13,19,26
	10	2,3,7,8,11,13,14,19,24,26
Euro28	5	1,2,4,6,9
	10	1,2,3,4,6,9,10,17,25,26

Table 1. Nodes in the EON where DCs are co-located in the different scenarios

Moreover, to set the popularities of the different CGs offered by the CDN, we rely on the zipf distribution, a frequently used approach in the literature to model the relative popularity of contents in the Internet [20] [21] [22] [23] [24]. In particular, sorting the CGs from the most to the least popular, the relative popularity CG in position $x \in \{1, \dots, |\mathcal{G}|\}$ becomes:

$$P(X = x) = \frac{x^{-\mu}}{\sum_{n=1}^{|\mathcal{G}|} 1/n^{\mu}} \quad (17)$$

where parameter μ is the so called skew factor of the zipf distribution. These popularities are then used to initialize α_g parameters accordingly. Besides, we model the anycast demand traffic taking into account the population in the proximity region to each EON node as follows. We assume one anycast demand initiated from each network node $v \in \mathcal{V}$ per CG $g \in \mathcal{G}$ with a traffic volume $T_{vg} = \alpha_g POP_v A$ Gb/s, where POP_v denotes the population in the proximity region to node v , and A is a factor used to adjust the offered traffic to the network. Specifically, if we want to offer a total traffic volume equal to T Gb/s, if POP denotes the global population in the network, $A = T/POP$, as $\sum_g \alpha_g = 1$. With these considerations, note that we are offering a total number of $|\mathcal{V}||\mathcal{G}|$ anycast demands to the network.

6.2. CPRMSA-PD performance validation

To validate the performance of CPRMSA-PD against optimal results obtained by solving model CPRMSA-ILP, we employ the medium-sized INT9 network, where CPRMSA-ILP incurs reasonable execution times. The two scenarios with 2 and 3 DCs (shown in Table 1) have been considered, assuming different replica server capacities. For instance, $n_r = \{250, 750\}$ indicates that replica servers at nodes 3 and 6 have capacities of 250cu and 750cu, respectively, while $n_r = \{250, 500, 250\}$ indicates that replica servers at nodes 3, 5 and 6 have respective capacities of 250cu, 500cu and 250cu. As for the CDN offered contents, a total content size of 1000cu has been considered, divided into 4 CGs of 250cu each, setting their popularity following the already introduced zipf distribution in eq. (17) with skew factor of 0.5. Therefore, according to our traffic model, we are offering $9 \cdot 4 = 36$ anycast demands to the network. The total offered traffic volume has been set to 4 Tbps, uniformly requested from all network nodes, since by now the population in the surrounding area of all network nodes is considered the same (i.e., every node initiates an anycast demand of identical volume per CG).

Regarding the specific values of the parameters involved in the CPRMSA-PD procedures, we have set parameters T_{coef} and τ in the SA-based candidate lightpath selection sub-problem (Figure 5) to 0.05 and 0.999, respectively, based on the findings in [15]. Moreover, after extensive tests in all considered scenarios (INT9, US26 and Euro28 networks) we have found that 100 global iterations (numOfGlobalIter) and 2500 iterations (maxIter) in the SA-based candidate lightpath selection procedure generally allow reaching good results with reasonably low execution times. In addition, to properly set the values for β and γ parameters involved in the CP and demand aggregation sub-problems, we have performed the following study. In four INT9 network scenarios with 3 DCs, we have initially fixed $\beta = 0.0$ and evaluated the number of slices used in the network (F) for different values of γ from 0.0 to 1.0. From these results, we take the most promising γ value and then use it to quantify F for different values of β , also from 0.0 to 1.0, identifying the best one as well. These results are depicted in Figure 7.

Looking at Figure 7 (left) we can see that the value of γ certainly has an impact on the number of slices used in the network, particularly in those scenarios with higher replica server capacity, and thus, CG replication. As a matter of fact, very limited replica server capacities do not allow for CG replication, which makes γ to have no effect as seen with $n_r = \{250, 500, 250\}$. In view of the results, we choose $\gamma = 0.2$ for CPRMSA-PD, and use it to observe CPRMSA-PD performance for different β values in Figure 7 (right). As seen, the value of β has an effect on the content placement performed by CPRMSA-PD that is eventually translated into different number of slices used. In general, it seems that small β values (i.e., higher load-balancing consideration in the content placement sub-problem) can provide the best performance of the heuristic. Therefore, we choose $\beta = 0.1$ for CPRMSA-PD. For practicality reasons, the presented values of all these parameters (T_{coef} , τ , numOfGlobalIter, maxIter, γ and β) will be kept fixed for all CPRMSA-PD executions throughout this section.

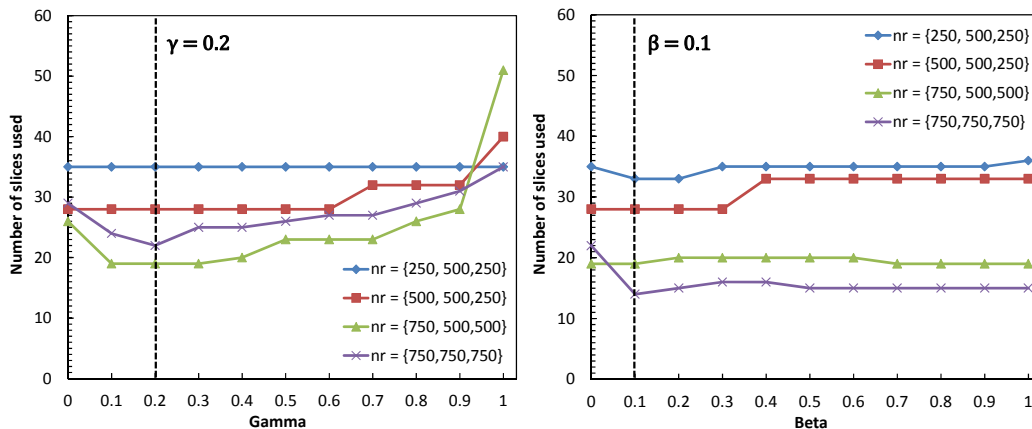


Figure 7. Number of slices used in the INT9 with 3 DCs using CPRMSA-PD: with parameter $\beta=0.0$ and for different values of parameter γ (left); with $\gamma = 0.2$ and for different values of β (right)

Once the validation scenario details have been presented, as well as the specific CPRMSA-PD parameter values have been selected, Table 2 compares the CPRMSA-PD performance against CPRMSA-ILP. Looking at the results, CPRMSA-PD yields in some scenarios exactly the same values of F as CPRMSA-ILP (i.e., 0% optimality gap), thus performing perfectly. Moreover, although the optimality gap can rise to 10-17% for some executions, absolute difference in the value of F remains quite low, up to +5 in the worst case. Focusing on the execution times, we

observe that they remain quite constant for CPRMSA-PD, around 1 minute in all tests. Conversely, they dramatically rise up for CPRMSA-ILP in some cases. In fact, we fixed the maximum execution time of CPLEX executions to 6 hours without reaching the target optimality gap that we set to 1%. As shown in the table, this time was exceeded in some cases, ending with the estimated optimality gap shown between parentheses. Again, although such CPLEX optimality gaps seem high, they correspond to the difference of only 1-2 slices in the optimized spectrum at most. Furthermore, CPLEX may take several additional hours, or even days to find these solutions.

$ \mathcal{R} $	n_r	CPRMSA-ILP		CPRMSA-PD		
		F	Total exec. time (s)	F	Opt. gap (%)	Total exec. time (s)
2	250,750	31	14852.06	31	0.0	75.80
	750,250	39	225.47	39	0.0	75.77
	500,500	31	>6h (6.2%)	35	12.9	78.33
	500,750	22	1474.55	23	4.5	70.69
	750,750	19	5713.69	21	10.1	69.06
	1000,1000	17	>6h (7.1%)	17	0.0	56.74
3	250,500,250	28	2943.30	33	17.8	76.66
	500,500,250	27	>6h (3.4%)	28	3.7	80.56
	750,500,500	18	1778.49	19	5.6	54.90
	750,750,750	12	6321.45	14	16.7	55.49

Table 2. CPRMSA-PD heuristic validation against the optimal CPRMSA-ILP: The number of slices used in the network F , optimality gap (in %) of CPRMSA-PD and total execution times are shown.

6.3. Numerical results in large network scenarios

Once validated the performance of CPRMSA-PD, we use the algorithm in this section to provide better understanding of the effects of the different parameters in a CDN over EON scenario. The results in this section will be extracted in the US26 and Euro28 reference backbone networks. According to CISCO traffic predictions in [5], Internet CDN traffic will be 322 Exabyte per year in 2015, that is, around 85 Tbps. US and Europe contribute to about 25% of the overall Internet traffic each. Therefore, we will assume an offered CDN traffic of 20 Tbps in both networks, following our traffic model while using real population estimates for each node region.

Table 3 shows the number of slices used in both US26 and Euro28 networks, as a function of the number of DCs, capacities of the replica servers in these DCs and number of CGs into which the total content size is divided. Specifically, 5 and 10 DCs are considered in each network, located as previously presented in Table 1. Moreover, 6 scenarios with different replica server capacities have been assumed in each case, which represent different $\lambda = \sum_r n_r / \sum_g h_g$ ratios, namely, the ratio of total replica server capacity in the network vs. total content size. Concerning the number of CGs, 10, 20 and 50 have been considered. Note that, with $|\mathcal{G}| = 50$, 1300 and 1400 anycast demands are offered to US26 and Euro28 networks, respectively. Regarding the CG popularity distribution function, the skew factor of 0.5 is kept as before.

$ \mathcal{R} $	n_r	λ	US26			Euro28		
			$ \mathcal{G} =10$	$ \mathcal{G} =20$	$ \mathcal{G} =50$	$ \mathcal{G} =10$	$ \mathcal{G} =20$	$ \mathcal{G} =50$
5	{200,200,200,200,200}	1.0	105	104	106	165	165	165

	{200,300,300,200,300}	1.3	82	83	78	145	151	160
	{400,300,300,200,300}	1.5	78	85	82	150	157	157
	{400,400,400,400,400}	2.0	64	66	65	139	137	141
	{600,600,600,600,600}	3.0	52	51	55	117	116	115
10	{100,100,100,100,100, 100,100,100,100,100}	1.0	123	122	119	170	170	171
	{200,100,100,200,100, 100,100,200,100,100}	1.3	80	91	102	152	152	146
	{200,100,200,200,200, 200,100,100,100,100}	1.5	73	74	74	129	133	124
	{200,200,200,200,200, 200,200,200,200,200}	2.0	56	57	63	86	84	92
	{300,300,300,300,300, 300,300,300,300,300}	3.0	36	37	37	59	63	62

Table 3. Number of slices used, F , in the US26 and Euro28 reference backbone networks with different number of DCs and replica server capacities. $|\mathcal{G}| = 10, 20$ and 50 are considered

Looking at the results, we clearly see that the number of slices used generally decreases inversely proportional to λ . This is due to the fact that the higher the capacities of the DC replica servers, the more the CGs can be replicated, thus placing them closer to the end users. As a result, the established lightpaths from the DCs to the users have to travel smaller distances (in km), which allows their assigned BV-TXPs to employ more spectrally-efficient modulation formats, hence requiring fewer spectrum slices eventually. Additionally, the more CGs we are allowed to place in the same DC the easier is to aggregate demands, as more anycast demands initiated from a certain network node can be delivered by that DC (it hosts the requested CGs). In contrast, if DC replica server capacities are scarce, CGs tend to be scattered among the network DCs, which makes demand aggregation from a specific DC and toward the same network node more difficult. This is translated into lower lightpath utilization (they tend to travel quite empty) and, thus, higher number of frequency slices used to carry the same amount of offered anycast demand. Further insight into these issues is provided later on upon discussing Figure 8. Another interesting detail to mention is the low effect of $|\mathcal{G}|$ on the number of slices used. This is reasonable as more CGs are translated into more anycast demands initiated from each network node (according to our traffic model). However, their traffic volume is lower, since the total offered load remains identical. The same happens with the size of the CGs with respect to the capacities of the DC replica servers, as the total content size does not change.

To complete the results in Table 3, we depict in Table 4 information about the average content placement time and total execution time of CPRMSA-PD needed to populate the previous table. Each of the time values shown is the average among the executions needed to solve all scenarios ($\lambda = 1.0, 1.3, 1.5, 2.0$ and 3.0) for that specific network, number of DCs and CGs. As can be observed, CPRMSA-PD succeeds in solving large problem instances with reasonable execution times, spending around 10-20 minutes in the US26 network, and up to 40 minutes in the Euro28 network largest problem instances. Concerning the CP time, we can see that it becomes negligible in all tests compared to the total CPRMSA-PD execution time. For example, even with 10 DCs and 50 CGs, average CP time stays in 12.43 and 16.58 seconds for US26 and Euro28 networks, where the CP ILP has to deal with 1300 and 1400 anycast demands, respectively (i.e., $|\mathcal{D}| = |\mathcal{V}||\mathcal{G}|$).

		US26		Euro28	
$ \mathcal{R} $	$ \mathcal{G} $	CP time (s)	Total exec. time (s)	CP time (s)	Total exec. time (s)
5	10	0.47	635.49	0.49	1464.45
	20	0.79	620.30	1.07	1558.43
	50	2.25	652.09	3.97	1712.20
10	10	1.16	838.13	2.48	2121.54
	20	2.68	923.98	5.60	2240.84
	50	12.43	1076.22	16.58	2370.88

Table 4. Average content placement and total execution time in seconds of CPRMSA-PD heuristic to obtain the results previously presented in Table 3.

In order to provide further insight into the actual lightpath allocation over the EON in the scenarios described in Table 3, we depict in Figure 8 the operational bit rate of the BV-TXP equipped in US26 and Euro28 networks with 5 DCs. Regarding the DC replica server capacities, $\lambda=1.5$ and 3.0 are analyzed (i.e., $n_r=\{400, 300, 300, 200, 300\}$ and $n_r=\{600,600,600,600,600\}$, respectively). Moreover, $|\mathcal{G}|=20$ has been chosen, which describes an intermediate scenario among the ones previously evaluated in Table 3.

Looking at the results, we can identify substantial differences in the required number of BV-TXPs, operational bit rates and modulation formats used when comparing $\lambda=1.5$ and $\lambda=3.0$ scenarios, both in US26 and EON networks, which justifies our initial findings from the results in Table 3. Specifically, with $\lambda=1.5$ CG replication among DCs is notably limited (i.e., only half of the CGs can be replicated in two DCs on average). According to our traffic model, however, all nodes initiate one anycast demand for every CG, which very frequently must be delivered from a distant DC (the only one hosting the desired CG). This requires BV-TXPs to make use of modulation formats with low spectral efficiency (like PM-BPSK, extensively used in the depicted $\lambda=1.5$ scenarios) in order to alleviate the transparent reach limitations. Besides, proper anycast demand aggregation becomes more difficult, as discussed before. In contrast, substantial CG replication is allowed in $\lambda=3.0$ scenarios (each CG can be replicated in 3 DCs on average). This effectively puts the CGs closer to the users demanding them. As a result, BV-TXPs can utilize more spectrally-efficient modulation formats (PM-8-QAM and PM-16-QAM instead of PM-BPSK) which require less spectral slices (as was previously identified in Table 3). Furthermore, as more CGs are hosted by DC replica servers, the aggregation of demands from a DC to a specific network node is more effective, which allows filling BV-TXPs capacities more efficiently. For example, in the US26 network with 5 DCs and $\lambda=1.5$, 116 lightpaths are established over the EON, 48 at 100 Gb/s, 36 at 200 Gb/s, 7 at 300 Gb/s and 25 at 400 Gb/s. Knowing that 1.9 of the offered 20 Tb/s are delivered from a DC located at the same node of the users initiating the demands, 18.1 Tb/s are carried over the network, which implies a 75% utilization of the operational BV-TXP capacity (39% of the maximum BV-TXP capacity). Conversely, in the same network with $\lambda=3.0$, only 83 lightpaths are needed, 29 at 100 Gb/s, 19 at 200 Gb/s, 7 at 300 Gb/s and 27 at 400 Gb/s. In this case 16.15 Tb/s are carried over the network, which represents an 82% utilization of the operational BV-TXP capacity (49% of the maximum BV-TXP capacity). Note that we omit the same analysis for US26 and Euro28 networks with 10 instead of 5 DCs, as they follow the same trends discussed here.

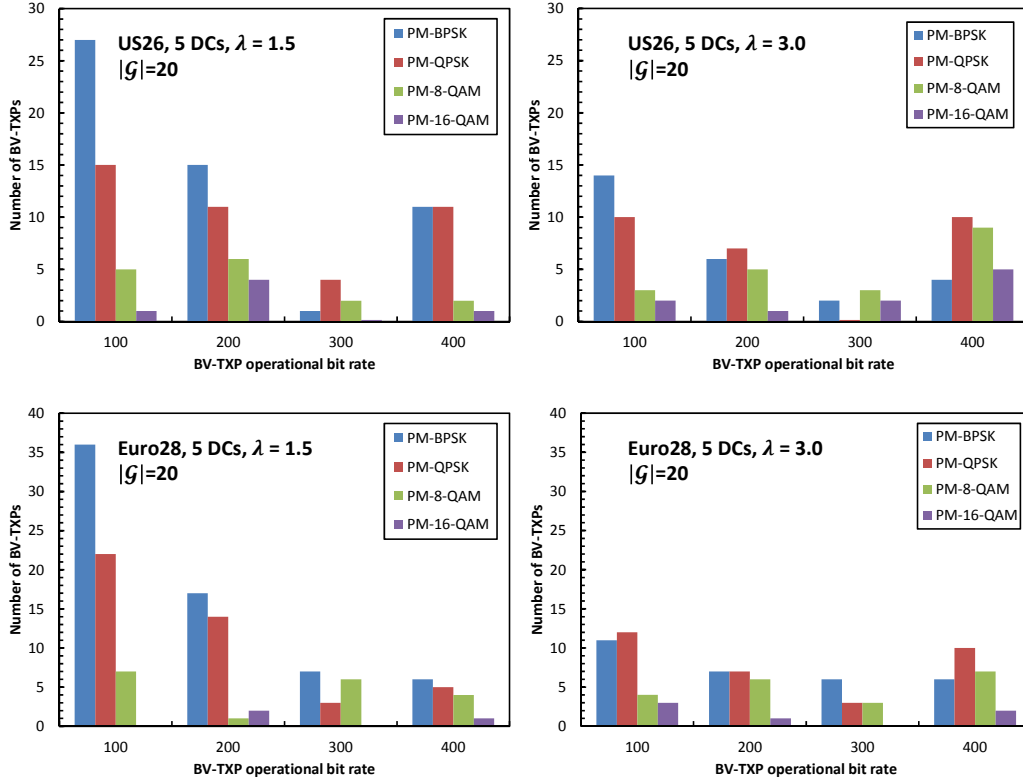


Figure 8. Operational bit rate and modulation format employed by the BV-TXP equipped in US26 (top row) and Euro28 (bottom row) networks with 5 DCs. Scenarios with $\lambda=1.5$ and $\lambda=3.0$ are considered.

Finally, we dedicate some more efforts on evaluating the impact of the CG popularity distribution function (i.e., the skew factor of the zipf distribution) on the utilization of the underlying EON. To this end, we take the Euro28 network with 5 and 10 DCs, and replica server capacities to have $\lambda=1.5$ and $\lambda=3.0$ (same n_r values as before). Regarding the number of CGs, $|\mathcal{G}|=10$ and $|\mathcal{G}|=50$ scenarios have been considered, which reflect the two extreme situations evaluated in previous tests. Taking all this into account, Table 5 shows the number of slices used in the EON as a function of the CG popularity distribution function skew factor (μ).

		Euro28 ($\lambda = 1.5$)		Euro28 ($\lambda = 3.0$)	
$ \mathcal{R} $	μ	$F(\mathcal{G} =10)$	$F(\mathcal{G} =50)$	$F(\mathcal{G} =10)$	$F(\mathcal{G} =50)$
5	0.3	147	153	112	117
	0.5	150	157	117	115
	0.7	150	164	118	116
	1.0	151	163	116	118
	1.2	157	163	118	118
10	0.3	135	135	71	62
	0.5	129	124	59	62
	0.7	117	123	61	64
	1.0	114	120	62	57
	1.2	119	116	59	60

Table 5. Number of slices used in the Euro28 reference backbone network with 5 and 10 DCs and replica server capacities to achieve $\lambda=1.5$ and $\lambda=3.0$, as a function of μ .

Looking at the table, we can identify some initially surprising yet comprehensive behaviors. While the number of frequency slices used seems to slightly increase with μ in the scenarios with 5 DCs, the opposite happens when the number of DCs rises to 10. Higher μ values imply that some CGs are highly demanded (their associated anycast demands are of high volume), whereas all the CGs remainder are barely popular (they attract anycast demands of low traffic volume). This tends to substantially increase the allocation of lightpaths working at low bit rates, particularly in the $\lambda=1.5$ scenarios where CGs are scattered across DCs and aggregation becomes difficult, which lowers the spectrum efficiency. In addition, some very high bit rate lightpaths are required to allocate the anycast demands for the more popular CGs that, although more spectrally-efficient (especially if advanced modulation formats can be employed), do not compensate for the large amount of underutilized lightpaths and results into using more frequency slices. This “bi-modal” lightpath setup is exemplified in Figure 9, where the configuration of lightpaths over the EON is shown when passing from $\mu=0.3$ to $\mu=1.2$ in the Euro28 with 5 DCs, $\lambda=3.0$ and $|\mathcal{G}|=10$.

The reason why the number of frequency slices used even decreases instead of increasing in the scenarios with 10 DCs is the following. In these cases, we have more than 1/3 of network nodes with a DC co-located. As said before, with high μ values a small amount of CGs attract most of the anycast demand traffic. Therefore, if only most popular CGs can be replicated in many DCs, most of the anycast demand volume requested from the nodes where DCs are co-located does not even have to enter the network, thus compensating for the lower spectral efficiency of the underutilized lightpaths that still have to be allocated. To show this in numbers, in the scenario with 5 DCs, $\lambda=3.0$ and $|\mathcal{G}|=10$, the total anycast demand traffic delivered from a co-located DC (does not enter the backbone EON) increases from 2.175 to 3.33 Tb/s when passing from $\mu=0.3$ to $\mu=1.2$. Nevertheless, in the same scenario but with 10 DCs, it increases from 2.65 to 5.1 Tb/s. This reduction of around 2.5 Tb/s in the traffic volume that the EON has to carry is really what leads to the reduction of frequency slices used, not the lower μ value itself. In general, such results highlight the importance of CG replication to foster demand aggregation, properly filling lightpath capacities and, thus, counteracting the negative effects of highly-skewed content popularity distributions.

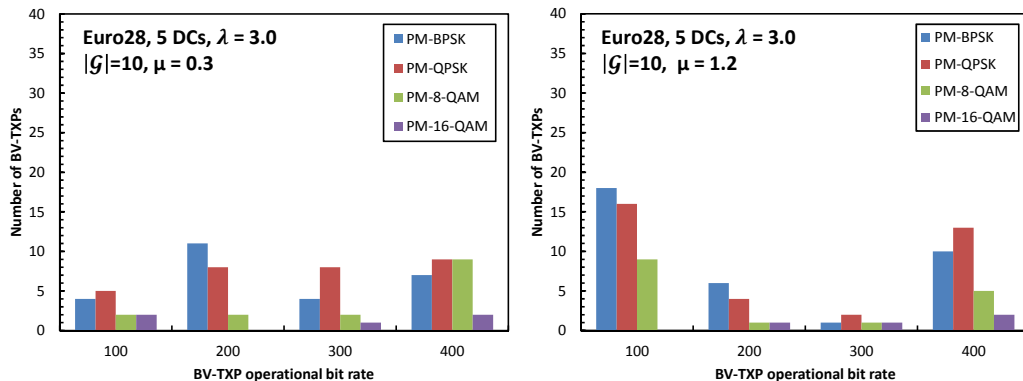


Figure 9. Operational bit rate and modulation format employed by the BV-TXP equipped in the Euro28 network when passing from CG popularity distribution function skew factor of 0.3 to 1.2

7. CONCLUDING REMARKS

This paper has addressed the joint Content Placement and lightpath Routing Modulation format and Spectrum Assignment (CPRMSA) problem in Content Distribution Network (CDN) over Elastic Optical Network (EON) scenarios. To solve it, an optimal ILP formulation (CPRMSA-ILP) as well as a heuristic based on problem decomposition (CPRMSA-PD) have been proposed. Once the performance of CPRMSA-PD has been successfully validated against the optimal CPRMSA-ILP model in medium-sized problem instances, we use it to analyze the effects that key parameters detected in the targeted scenario (number of DCs, replica server capacity, number of content groups, and content group popularity distribution function) can have on the utilization of the underlying EON resources. From the obtained results, sufficient capacity to allow content replication in the DCs yields dramatic reductions in terms of frequency slices used, allowing for superior anycast demand aggregation, together with the usage of advanced modulation formats (PM-8-QAM, PM-16-QAM instead of PM-BPSK or PM-QPSK). Furthermore, while no substantial effects of the number of content groups on the required number of frequency slices are observed, we identify potential EON underutilization in highly-skewed content group popularity distributions if very limited content replication prevents effective demand aggregation.

It should be noted that the addressed optimization problem is static and assumes that popularity of Content Groups is fixed. However, the proposed approach could be also applied in such scenarios that take into account evolving popularity trends. For example, the proposed static optimization process could be repeated periodically (or triggered upon significant CG popularity changes) using the new data at the moment. Next, reconfiguration to the new optimal setup would follow. Alternatively, the proposed static optimization process could only account for CGs whose popularity is expected to remain quite constant over time. Then, for those CGs whose popularity is expected to show higher variability over time, as well as for the associated anycast demands, lighter dynamic CP and lightpath RMSA heuristics and procedures could be employed. Both applications are identified as possible research lines and left for future work at this point.

ACKNOWLEDGEMENTS

This work has been partially funded by the Spanish National project SUNSET (TEC2014-59583-C2-1-R/TEC2014-59583-C2-2-R) and the *Departament d'Universitats, Recerca i Societat de la Informació de la Generalitat de Catalunya* under contract 2014SGR-1427. Moreover, the work of Krzysztof Walkowiak was supported by the Polish National Science Centre (NCN) under Grant DEC-2012/07/B/ST7/01215.

REFERENCES

- [1]. B. Ahlgren, C. Dannewitz, C. Imbrenda, D. Kutscher, and B. Ohlman, "A survey of information-centric networking", *IEEE Communications Magazine*, vol. 50, no. 7, pp. 26-36, 2012.
- [2]. C. Jaeyoung, H. Jinyoung, C. Eunsang, T. Kwon, and C. Yanghee, "A Survey on content-oriented networking for efficient content delivery", *IEEE Communications Magazine*, vol. 49, no. 3, pp.121-127, 2011.

- [3]. C. Develder, M. De Leenheer, B. Dhoedt, M. Pickavet, D. Colle, F. De Turck, and P. Demeester, "Optical networks for grid and cloud computing applications", Proceedings of the IEEE, vol. 100, no. 5, pp. 1149-1167, 2012.
- [4]. M. Klinkowski and K. Walkowiak, "On advantages of elastic optical networks for provisioning of cloud computing traffic", IEEE Network Magazine, vol. 27, no. 6, pp. 44-51, 2013.
- [5]. Cisco, "Cisco Visual Networking Index: Forecast and Methodology, 2014–2019", Cisco White Paper, 2015. http://www.cisco.com/c/en/us/solutions/collateral/service-provider/ip-ngn-ip-next-generation-network/white_paper_c11-481360.html
- [6]. O. Gerstel, M. Jinno, and A. Lord, "Elastic optical networking: A new dawn for the optical layer?", IEEE Communications Magazine, vol. 50, no. 2, 2012.
- [7]. R. Goścień, K. Walkowiak, and M. Klinkowski, "Distance-adaptive transmission in cloud-ready elastic optical networks", IEEE/OSA Journal of Optical Communications and Networking, vol. 6, no. 10, pp. 816-828, 2014.
- [8]. M. Klinkowski, and K. Walkowiak, "Routing and spectrum assignment in spectrum sliced elastic optical path network", IEEE Communications Letters, vol. 15, no. 8, 2011.
- [9]. K. Christodoulopoulos, I. Tomkos, and E. Varvarigos, "Elastic bandwidth allocation in flexible OFDM based optical networks", IEEE Journal of Lightwave Technology, vol. 29, no. 9, pp. 1354–1366, 2011.
- [10]. Y. Wang, X. Cao, and Y. Pan, "A study of the routing and spectrum allocation in spectrum-sliced elastic optical path networks", in Proceedings of IEEE INFOCOM, Shanghai, China, 2011.
- [11]. S. Talebi, E. Bampis, G. Lucarelli, I. Katib, and G. N. Rouskas, "Spectrum assignment in optical networks: A multiprocessor scheduling perspective", Journal of Optical Communications and Networking, vol. 6, no. 8, pp. 754–763, 2014.
- [12]. L. Velasco, M. Klinkowski, M. Ruiz, and J. Comellas, "Modeling the routing and spectrum allocation problem for flexgrid optical networks", Photonic Network Communications, vol. 24, no. 3, pp. 177–186, 2012.
- [13]. M. Żotkiewicz, M. Pióro, M. Ruiz, M. Klinkowski, and L. Velasco, "Optimization models for flexgrid elastic optical networks", in Proceedings of IEEE ICTON, Cartagena, Spain, 2013.
- [14]. M. Klinkowski, "An evolutionary algorithm approach for dedicated path protection problem in elastic optical networks", Cybernetics and Systems, vol. 44, no. 6-7, pp. 589–605, 2013.
- [15]. M. Aibin and K. Walkowiak, "Simulated annealing algorithm for optimization of elastic optical networks with unicast and anycast traffic", in Proceedings of IEEE ICTON, Graz, Austria, 2014.
- [16]. R. Goścień, K. Walkowiak, and M. Klinkowski, "Tabu search algorithm for routing, modulation and spectrum allocation in elastic optical network with anycast and unicast traffic", Computer Networks, vol. 79, pp. 148–165, 2015.
- [17]. K. Walkowiak, M. Klinkowski, B. Rabięga, and R. Goścień, "Routing and spectrum allocation algorithms for elastic optical networks with dedicated path protection", Optical Switching and Networking, vol. 13, pp. 63–75, 2014.
- [18]. K. Walkowiak and M. Klinkowski, "Joint anycast and unicast routing for elastic optical networks: Modeling and optimization", in Proceedings of IEEE ICC, Budapest, Hungary, 2013.

- [19]. R. Gościęć, K. Walkowiak, and M. Klinkowski, "Joint Anycast and Unicast Routing and Spectrum Allocation with Dedicated Path Protection in Elastic Optical Networks", in Proceedings of the 10th International Conference on Design of Reliable Communication Networks (DRCN 2014), Ghent, Belgium, 2014.
- [20]. A. Abhari and M. Soraya., "Workload generation for youtube. Multimedia Tools and Applications", vol. 46, no. 1, pp. 91–118, 2010.
- [21]. P. Gill, M. Arlitt, Z. Li, and A. Mahanti, "Youtube traffic characterization: A view from the edge", In Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement, IMC'07, pp. 15–28, 2007.
- [22]. L. Breslau, Pei Cao, Li Fan, G. Phillips, and S. Shenker, "Web caching and zipf-like distributions: evidence and implications", In Proceedings of Annual Joint Conference of the IEEE Computer and Communications (INFOCOM 1999), vol. 1, pp. 126–134, 1999.
- [23]. K. Cho, M. Lee, K. Park, T.T. Kwon, Y. Choi, and S. Pack, "Wave: Popularity-based and collaborative in-network caching for content-oriented networks", In Proceedings of the IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS 2012), pp. 316–321, 2012.
- [24]. T. Janaszka, D. Bursztynowski, and M. Dzida, "On popularity-based load balancing in content networks", In Proceedings of the 24th International Teletraffic Congress (ITC 2012), pp. 1–8, Sept 2012.
- [25]. G. Held, "A Practical Guide to Content Delivery Networks", Auerbach Publications, Boston, MA, USA, 2005.
- [26]. M. Hofmann, L. Beaumont, "Content Networking: Architecture, Protocols, and Practice", Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2005.
- [27]. Jiang Xiaoke, Bi Jun, Nan Guoshun, Li Zhaogeng, "A survey on Information-centric Networking: Rationales, designs and debates," in Communications, China , vol.12, no.7, pp.1-12, July 2015.
- [28]. T. Bektas, O. Oguz, and I. Ouveysi, "Designing cost-effective content distribution networks", Computers & Operations Research, vol. 34, no. 8, pp. 2436-2449, 2007.
- [29]. J. He, X. Zhao, and B. Zhao, "A fast, simple and near-optimal content placement scheme for a large-scale VoD system", in Proceedings of IEEE ICCS, 2012.
- [30]. A. Araldo, M. Mangili, F. Martignon, D. Rossi, Cost-aware caching: Optimizing cache provisioning and object placement in ICN, In Proceedings of the IEEE Global Communications Conference (GLOBECOM 2014), pp. 1108-1113, 2014.
- [31]. W. Fang, Z. Wang, J. Lloret, D. Zhang, Z. Yang, "Optimising data placement and traffic routing for energy saving in Backbone Networks", Transactions on Emerging Telecommunications Technologies, Vol. 25, No. 9, pp. 2161-3915, 2014.
- [32]. S. Ferdousi, F. Dikhiyik, M.F. Habib, B. Mukherjee, "Disaster-aware data-center and content placement in cloud networks," in Proceedings of IEEE International Conference on Advanced Networks and Telecommunications Systems (ANTS) 2013, 2013.
- [33]. S. Ferdousi, F. Dikhiyik, M. F. Habib, M. Tornatore, and B. Mukherjee, "Disaster-Aware Dynamic Content Placement in Optical Cloud Networks", In Proceedings of OFC 2014, San Francisco (CA) USA, 2014.
- [34]. Yonggong Wang, Zhenyu Li, G. Tyson, S. Uhlig, Gaogang Xie, "Optimal cache allocation for Content-Centric Networking", In Proceedings of IEEE ICNP 2013, Gottingen, Germany, October 2013.

- [35]. M. Mangili, F. Martignon, A. Capone, "A Comparative Study of Content-Centric and Content-Distribution Networks: Performance and Bounds", in Proceedings of IEEE Globecom 2013, Atlanta, GA, USA, December 2013.
- [36]. M. F. Habib, M. Tornatore, M. Leenheer, F. Dikbiyik, and B. Mukherjee, "Design of Disaster-Resilient Optical Data Center Networks", IEEE Journal of Lightwave Technology, vol. 30, no. 16, pp. 2563-2573, 2012.
- [37]. S. Ferdousi, F. Dikbiyik, M.F. Habib, M. Tornatore, B. Mukherjee, "Disaster-aware datacenter placement and dynamic content management in cloud networks," in Optical Communications and Networking, IEEE/OSA Journal of , vol.7, no.7, pp.681-694, July 2015.
- [38]. El-Ghazali Talbi, "Metaheuristics: From Design to Implementation", Wiley Ed., 2009.