

Manycast, Anycast, and Replica Placement in Optical Inter-Datacenter Networks

Ajmal Muhammad¹, Nina Skorin-Kapov², Marija Furdek¹

¹ Royal Institute of Technology (KTH), Stockholm, Sweden, Email: {ajmalmu, marifur}@kth.se

² University Centre of Defense, San Javier Air Force Base, Santiago de la Ribera, Spain, Email: nina.skorinkapov@tud.upct.es

Abstract—The expanding adoption of cloud-based services in recent years puts stringent requirements on datacenters (DCs) and their interconnection networks. Optical inter-datacenter networks represent the only viable option for satisfying the huge bandwidth required to replicate and update content for cloud-based services across geographically dispersed datacenters. In addition to content replication and synchronization, optical inter-datacenter networks must also support communication between datacenters and end-users. The resulting new traffic patterns and the enormous traffic volumes call for new capacity-efficient approaches for inter-datacenter network design that incorporate both transport and datacenter resource planning.

This paper introduces an integrated approach to optimally place content replicas across DCs by concurrently solving the routing and wavelength assignment (RWA) problem for both inter-DC content replication and synchronization traffic following the manycast routing paradigm, and end-user driven user-to-DC communication following the anycast routing paradigm, with the objective to reduce the overall network capacity usage. To attain this goal, the Manycast, Anycast and Replica Placement (MARP) problem is formulated as an integer linear program (ILP), to find optimal solutions for smaller problem instances. Due to the problem complexity, a scalable and efficient heuristic algorithm is developed to solve larger network scenarios. Simulation results demonstrate that the proposed integrated MARP strategy can significantly reduce the network capacity usage when compared to benchmarking replica placement and RWA schemes aimed at minimizing the resources consumed by either of the two types of traffic independently.

Index Terms— Content placement, inter-datacenter networks, manycast routing, anycast routing.

I. INTRODUCTION

As demands for cloud services proliferate, cloud Content Service Providers (CSPs) such as Amazon, Google, Facebook, etc, increasingly create, store and share massive amounts of content. Generally, content is replicated across multiple geographically dispersed datacenters (DCs) interconnected via ultra-high capacity wavelength division multiplexing (WDM)-based inter-DC networks, either owned or leased by the CSP [1]. Such distribution offers high service availability and enables end-users to connect to the most convenient (e.g., the closest) DC, thereby decreasing transit latency and/or the amount of network resources needed to support the service requests. In the context of network planning, cloud services give rise to new traffic patterns and huge traffic volumes, requiring tailored resource-efficient solutions [2].

The traffic associated with inter-DC networks can be classified into two broad categories. The first encompasses inter-DC traffic which is related to content replication and synchronization/updates between multiple DCs [3]. Routing paradigms associated to this type of traffic usually employ

unicast communication between DCs or multicast communication between a set of DCs that host content replicas [4]. The second type of traffic comprises end-user-driven communication where users access cloud content and services, typically by applying the anycast routing paradigm [1]. In anycast routing, a user demand is served from any one of multiple DCs which host the specified content/service. In other words, if replicas of a certain content are stored at multiple DCs, the user can connect to any one of them. This modifies the assumptions used by classic routing strategies because the traffic matrix is unknown, i.e., the destination node is not specified but can be any node within a subset. Both of the aforementioned traffic categories imply high bandwidth requirements and their provisioning should be optimized to achieve resource-efficient network operation.

The bandwidth strain induced by inter-DC content replication and synchronization traffic increases with the number of replicas and the distance between the hosting DCs. On the other hand, a large number of geographically distributed replicas reduces end-user-driven traffic bandwidth requirements (and transit latency) since the users can generally access a closer DC. Thus, efficient replica placement and resource allocation requires the development of advanced network planning strategies considering both traffic types.

Motivated by this observation, in this paper we solve the Manycast, Anycast and Replica Placement (MARP) problem in WDM-based inter-DC networks. Given an inter-DC network topology and DC locations, a set of contents to be replicated, a range of required replicas for each content and a set of user demands for specific content, our goal is to determine the number and placement of replicas for each content, as well as the associated routing and wavelength assignment for both traffic categories, i.e., inter-DC and user-to-DC traffic, with the objective of minimizing the total network capacity usage. User-driven demands are assumed to follow the anycast routing paradigm where users can establish an optical connection, or lightpath, to any DC hosting a replica of the desired content. Given a replica placement, the synchronization/update inter-DC traffic is assumed to employ the multicast routing paradigm where all DCs hosting content replicas are interconnected by a light-tree rooted at the main data center. Since we assume the replica placement is unknown, the manycast routing scheme [5] is applied, which interconnects only a subset of nodes (i.e., DCs) from the possible destination set. The end nodes of the obtained manycast tree then determine the replica placement. We formulate the Manycast, Anycast and Replica Placement (MARP) problem as an Integer Linear Program (ILP) with the objective to minimize the total network capacity usage. Preliminary results of the ILP were

presented in [6]. We extend upon our previous work here by developing an effective and scalable heuristic algorithm for larger instances and present new simulation results comparing with benchmarking strategies which consider the two traffic types independently.

The remainder of this paper is structured as follows. An overview of related work is presented in Section II. Section III presents the ILP formulation for the MARP problem while Section IV describes the proposed heuristic approach. Numerical results are analyzed in Section V, and concluding remarks are given in Section VI.

II. RELATED WORK

The establishment of anycast and/or unicast demands in the context of inter-DC networks, as well as replica placement schemes, have been individually studied in the literature. General concepts of anycast and anycast routing and wavelength assignment approaches for optical networks have been investigated in [5], [7]–[12]. In [7], the anycast communication paradigm is employed in multi-layer networks to improve WDM network response to IP traffic variations. In [8], the authors perform physical-layer impairment-aware anycast routing to decrease the blocking probability of requests caused by impairments. In [9], the authors study optical anycast in networks equipped with wavelength-converters while trying to minimize the total number of necessary converters. An optimization framework for the static anycast routing and wavelength assignment problem with the objective to minimize the highest used wavelength in the network is presented in [5], while the approach in [11] aims at minimizing the total amount of resources used by the anycast demands. The authors in [10] present energy-efficient schemes for anycast implemented as an overlay on the optical layer that lacks the ability to switch/split optical signals to more than one destination.

In addition to the above mentioned general approaches for anycast and anycast communication, some of which may be applied to datacenter networks, tailored anycast and anycast routing strategies for end-user and/or inter-DC traffic under a given replica placement have also been proposed, typically aimed at improving resource usage efficiency, reducing the power consumption or enhancing network survivability in the presence of failures. The savings in network resource usage enabled by multicast routing of inter-DC traffic compared to unicast connections between DCs are studied in [13]. Anycast routing of static end-user traffic demands to a given DC set using a minimum amount of network resources is explored for both working and backup paths to provide resilience against link and node (DC) failures in [14]. Approaches for single-path and multi-path routing of dynamic anycast demands aimed at minimizing demand blocking can be found in [15]. In [16], the authors investigate the impact of energy-aware anycast routing and wavelength assignment to energy savings in optical datacenter networks.

The problem of placing content replicas over geographically distributed datacenters and its impact on various network parameters, such as cost, power consumption, QoS or security has been investigated in [17]–[25]. Replica placement approaches aimed at minimizing the bandwidth and storage cost for content distribution networks can be found in [17], [19]. The work in [20] focuses on replica placement

to reduce the total power consumption of the network, while the study in [21] aims at enhancing content availability in datacenter networks. In [22], the authors assess the network power savings attainable by storing the frequently accessed contents in the proximity of their users. The authors in [23] present an optimization framework for datacenter site selection that takes into account the CAPEX and OPEX costs of possible locations, response time to user requests, service availability, and CO2 emission, in order to derive strategic tradeoffs between these parameters. The study in [24] addresses the problem of DC selection for joint minimization of non-renewable energy consumption and the number of datacenters. Efficient failure-resilient replica placement strategies can be found in [25] for protecting inter-DC traffic against a critical node (DC) failure during transmission from the main datacenter to replica-hosting DCs.

Some studies have also coupled the content replica placement for datacenter networks with the routing of the end-user demands considering different objectives. For example, the authors in [26] study the problem of replica placement and anycast routing of end-user demands to minimize the total network resource usage in elastic optical networks. Similarly, the authors in [27] combine content placement with anycast routing for survivable cloud service provisioning to minimize the total cost of the network. The study in [12] investigates joint and subsequent approaches for datacenter selection and the routing of end-user traffic demands. The authors in [1] consider content placement, the routing of end-user requests, and protection of both paths (for end-user traffic) and contents against a single disaster failure simultaneously, with the objective to reduce the network resource usage. The work in [28] proposes a disaster-aware service provisioning scheme that employs anycast routing to connect end-users to a subset of datacenters for survivability purposes and also to distribute the required bandwidth for requested services over multiple paths destined to DCs hosting the service/content.

The study in [4] models the replica placement as soft-capacitated connected facility location problem which optimizes both user-driven traffic and inter-DC replication/update traffic but only for a single content and without considering wavelength allocation. Also, there is no requirement on the number of replicas which may not be practical in the context of survivability considering that a lower number or more geographically concentrated user demands would lead to solutions with few or no replicas. To the best of the authors' knowledge, no prior studies have taken into account the impact of both types of traffic on replica placement for a set of contents to reduce the overall resource usage in optical datacenter networks. Considering the fact that inter-datacenter traffic accounts for up to 45% of the total traffic in the backbone network [29] and constitutes 9% share of the global datacenter traffic with an estimated annual growth rate of 31.9% [2], it is important to incorporate the traffic related to content replication and synchronization into network planning strategies. To fill this gap, we propose a planning strategy for WDM-based optical inter-DC networks that jointly considers the routing and wavelength assignment of end-user demands and update/synchronization traffic when placing replicas of a set of contents, with the overall objective of decreasing the usage of network resources. The joint approach of considering both traffic types simultaneously leads to enhanced solutions in comparison to solving the individual

problems subsequently. The planning strategy presented in this paper does not consider physical layer impairments. However, the model can easily be extended to networks with physical layer impairments by modifying the cost function described in Section IV.A.

III. OPTIMIZATION MODEL FOR THE MARP PROBLEM: MARP-ILP

An optical inter-datacenter mesh core network is considered where each DC is assumed to be built in the proximity of one core node, and connected to it by access or metro links, not reflected in this study. We assume that a bidirectional lightpath is established for each anycast demand for a content. Note that the network nodes are assumed to represent core nodes which collect and aggregate several individual user requests into a set of anycast demands for specific content and employ traffic grooming to efficiently utilize the wavelength capacity. In addition to these anycast demands, a light-tree is established for each content interconnecting all the nodes corresponding to DCs hosting replicas of that content (i.e., multicast demands). Although we consider a light-tree implementation in this work, the approach could easily be extended to consider multicast trees supported by a set of lightpaths.

The objective of the Multicast, Anycast and Replica Placement (MARP) problem considered in this paper is to establish a set of lightpaths and light-trees for the DC-to-DC and user-to-DC traffic, respectively, and perform replica placement such that the overall network resource consumption is minimized. This implies solving the routing and wavelength assignment (RWA) problem for both end-user demands and update/synchronization traffic, concurrently, by applying the anycast and multicast routing paradigms, while the obtained solution yields a replica placement. We model the MARP problem as an Integer Linear Program (ILP) as follows.

Input parameters

- $\mathcal{G}(\mathcal{V}, \mathcal{E})$: a directed graph where \mathcal{V} is the set of vertices that represent the network nodes, \mathcal{E} is the set of edges that represent the network links;
- \mathcal{C} : set of contents;
- $\mathcal{D} \in \mathcal{V}$: set of datacenter locations;
- $k_{\min}(k_{\max})$: minimum (maximum) number of replicas required for each content;
- \mathcal{S} : maximum number of replicas a datacenter can store;
- \mathcal{M} : set of multicast demands (d, c) , where $d \in \mathcal{D}$ is the root/main datacenter for content $c \in \mathcal{C}$;
- \mathcal{A} : set of anycast demands (s, c) , where $s \in \mathcal{V}$ is the source node and $c \in \mathcal{C}$ is the requested content;
- \mathcal{W} : set of available wavelengths;
- $\Upsilon = \{K_1, K_2, \dots, K_{|\mathcal{D}|}\}$: set of constants with values such that $\sum_{i=1}^{\hat{i}-1} K_i \ll K_{\hat{i}}$, used for assigning a unique identifier to each node of \mathcal{D} ;

Variables

- $y_{(i,j)}^{(d,c),w} \in \{0, 1\}$ - equal to 1 if $w \in \mathcal{W}$ is used on link $(i, j) \in \mathcal{E}$ for multicast demand (d, c) and 0 otherwise;
- $U_i^{(d,c)} \in \mathbb{Z}$ - denote the order of vertices included in the route tree for multicast demand (d, c) ;
- $\bar{y}_{(i,j)}^{(s,c),w} \in \{0, 1\}$ - equal to 1 if $w \in \mathcal{W}$ is used on link $(i, j) \in \mathcal{E}$ for anycast demand (s, c) and 0 otherwise;

- $N_d^c \in \{0, 1\}$ - equal to 1 if $d \in \mathcal{D}$ is selected as one of the location for c and 0 otherwise;
- $k^c \in \mathbb{Z}$ - the number of datacenters selected for hosting replicas of content c ;
- $\bar{A}_d^{(s,c)} \in \{0, 1\}$ - equal to 1 if d is selected as the destination node for anycast demand (s, c) and 0 otherwise;
- $\Omega^{(d,c),w} \in \{0, 1\}$ - equal to 1 if w is used by (d, c) and 0 otherwise;
- $\bar{\Omega}^{(s,c),w} \in \{0, 1\}$ - equal to 1 if w is used by (s, c) and 0 otherwise;

Objective function

$$\text{Minimize } \sum_{(d,c) \in \mathcal{M}} \sum_{(i,j) \in \mathcal{E}} \sum_{w \in \mathcal{W}} y_{(i,j)}^{(d,c),w} + \sum_{(s,c) \in \mathcal{A}} \sum_{(i,j) \in \mathcal{E}} \sum_{w \in \mathcal{W}} \bar{y}_{(i,j)}^{(s,c),w} \quad (1)$$

Constraints

$$\sum_{d \in \mathcal{D}} N_d^c = k^c, \quad \forall c \in \mathcal{C} \quad (2)$$

$$k_{\min} \leq k^c \leq k_{\max}, \quad \forall c \in \mathcal{C} \quad (3)$$

$$\sum_{c \in \mathcal{C}} N_d^c \leq \mathcal{S}, \quad \forall d \in \mathcal{D} \quad (4)$$

$$\sum_{j:(d,j) \in \mathcal{E}} \sum_{w \in \mathcal{W}} y_{(d,j)}^{(d,c),w} \geq 1, \quad \forall (d, c) \in \mathcal{M} \quad (5)$$

$$\sum_{i:(i,d) \in \mathcal{E}} \sum_{w \in \mathcal{W}} y_{(i,d)}^{(d,c),w} = 0, \quad \forall (d, c) \in \mathcal{M} \quad (6)$$

$$\sum_{d \in \mathcal{D}} \left(\sum_{i:(i,d) \in \mathcal{E}} \sum_{w \in \mathcal{W}} K_d \cdot y_{(i,d)}^{(d,c),w} - K_d \cdot N_d^c \right) = 0, \quad \forall (d, c) \in \mathcal{M} \quad (7)$$

$$\sum_{j:(i,j) \in \mathcal{E}} y_{(i,j)}^{(d,c),w} - |\mathcal{V}| \cdot \sum_{j:(j,i) \in \mathcal{E}} y_{(j,i)}^{(d,c),w} \leq 0, \quad \forall (d, c) \in \mathcal{M}, \forall w \in \mathcal{W}, \forall i \in \mathcal{V} \setminus \mathcal{D} \quad (8)$$

$$\sum_{j:(j,i) \in \mathcal{E}} y_{(j,i)}^{(d,c),w} - \sum_{j:(i,j) \in \mathcal{E}} y_{(i,j)}^{(d,c),w} \leq 0, \quad \forall (d, c) \in \mathcal{M}, \forall w \in \mathcal{W}, \forall i \in \mathcal{V} \setminus \mathcal{D} \quad (9)$$

$$U_i^{(d,c)} - U_j^{(d,c)} + |\mathcal{V}| \cdot y_{(i,j)}^{(d,c),w} \leq |\mathcal{V}| - 1, \quad \forall (d, c) \in \mathcal{M}, \forall w \in \mathcal{W}, \forall (i, j) \in \mathcal{E} \quad (10)$$

$$\sum_{w \in \mathcal{W}} \Omega^{(d,c),w} = 1, \quad \forall (d, c) \in \mathcal{M} \quad (11)$$

$$y_{(i,j)}^{(d,c),w} + y_{(j,i)}^{(d,c),w} \leq \Omega^{(d,c),w}, \quad \forall (d, c) \in \mathcal{M}, \forall w \in \mathcal{W}, \forall (i, j) : |j| > |i| \quad (12)$$

$$\sum_{i:(i,m) \in \mathcal{E}} \sum_{w \in \mathcal{W}} \bar{y}_{(i,m)}^{(s,c),w} - \sum_{j:(m,j) \in \mathcal{E}} \sum_{w \in \mathcal{W}} \bar{y}_{(m,j)}^{(s,c),w} = \begin{cases} -1, m = s \\ \bar{A}_d^{(s,c)}, \quad \forall d \in \mathcal{D}, \forall (s, c) \in \mathcal{A} \\ 0, \text{ otherwise} \end{cases} \quad (13)$$

$$\sum_{d \in \mathcal{D}} \bar{A}_d^{(s,c)} = 1, \quad \forall (s, c) \in \mathcal{A} \quad (14)$$

$$\bar{A}_d^{(s,c)} \leq N_d^c, \quad \forall (s,c) \in A, \forall c \in C, \forall d \in D \quad (15)$$

$$\sum_{w \in W} \bar{\Omega}^{(s,c),w} = 1, \quad \forall (s,c) \in A \quad (16)$$

$$\bar{y}_{(i,j)}^{(s,c),w} + \bar{y}_{(j,i)}^{(s,c),w} \leq \bar{\Omega}^{(s,c),w}, \quad (17)$$

$$\forall (s,c) \in A, \forall w \in W, \forall (i,j) : |j| > |i|$$

$$\sum_{(d,c) \in M} \bar{y}_{(i,j)}^{(d,c),w} + \sum_{(s,c) \in A} \bar{y}_{(i,j)}^{(s,c),w} \leq 1, \quad \forall (i,j) \in \mathcal{E}, \forall w \in W \quad (18)$$

The objective (1) is to minimize the number of wavelength-links required to establish the requested multicast light-trees and anycast lightpaths. Constraints (2,3) select the number of replicas to be hosted in DC locations for each content. Constraint (4) ensures that the number of replicas hosted in a DC for all the contents does not exceed the DC storage capacity. Constraints (5-10) are used to build each multicast light-tree. Constraint (5) specifies that there must be at least one wavelength associated with content c outgoing from the main DC node. Constraint (6) prohibits any incoming wavelengths to the main DC location. Constraint (7) ensures that the destination nodes of each tree host the associated content replicas. Constraint (8) specifies that each node other than the main DC can have outgoing wavelengths only if it has incoming wavelengths. Constraint (9) ensures that nodes which are not in the set of possible DC locations and have an incoming wavelength must have at least one outgoing wavelength. Note that the last two constraints (i.e., constraints (8,9)) differ from the flow conservation constraints, since the number of outgoing wavelength links for intermediate nodes in a light-tree can be greater than the number of input wavelength links due to branching of the tree. Constraint (10) prevents the formation of loops in established paths. Constraints (11,12) enforce the wavelength continuity constraint for multicast light-trees.

Constraints (13-15) are used to establish the anycast lightpaths. Constraint (13) enforces flow conservation for each lightpath by making sure that the number of outgoing and incoming wavelength links is equal for every node along its path, except for the source node of the demand and the selected destination DC node. Constraint (14) ensures that only one DC is assigned per anycast demand. Constraint (15) guarantees that a DC can be selected for a demand if and only if the requested content is replicated at that DC. Constraints (16,17) enforce the wavelength continuity constraint for the anycast lightpaths. Finally, constraint (18) enforces the wavelength clash constraint by making sure that each wavelength on a certain network link is allocated to at most one demand.

A. The Complexity of MARP-ILP

The multicast routing problem itself (also known as the k -Steiner tree problem [11]), as well as anycast routing and wavelength assignment, are NP-complete [30], [31]. Thus, the proposed ILP model for the combined MARP problem is tractable only for relatively small problem instances. For larger problems, obtaining practical solutions in reasonable time calls for heuristic approaches.

To get insight into the complexity of the proposed ILP formulation for MARP, we calculate the number of variables and constraints as a function of the set of contents (C), the set of datacenter locations (D), the set of multicast demands (M), the set of anycast demands (A), and the set of available wavelengths (W). The set of datacenter locations (D), the set of multicast demands (M), and the set of anycast demands (A) can be upper bounded by $O[|\mathcal{V}|]$, $O[|W||\mathcal{E}|]$, and $O[|W||\mathcal{E}|]$, respectively. Moreover, if we assume an asymptotic number of physical links, i.e., a fully connected physical topology, then $|\mathcal{E}| \approx |\mathcal{V}|^2$. Consequently, both the number of variables and the number of constraints grow in order of $O[|W|^2|\mathcal{V}|^4]$.

IV. HEURISTIC ALGORITHM FOR THE MARP PROBLEM: MARP-H

This section presents a scalable and efficient heuristic algorithm for the MARP problem. The algorithm, denoted as MARP-Heuristic (MARP-H), iteratively constructs the multicast trees for each content, and the associated anycast lightpaths, by adding the most beneficial DC to the replica location set in each step according to a cost factor, and adaptively updating the routes of the remaining anycast and multicast demands. First-fit wavelength assignment is then used to assign wavelengths to the established light-trees and lightpaths.

A. MARP-H Description

The pseudocode of the routing subproblem of MARP-H is shown in Fig. 1. First, the contents in C are sorted in descending order of the number of end-user demands for each of them. Then, for each content c , a multicast tree is constructed as follows. It is first rooted at the main DC and then expanded by iteratively adding DCs and merging paths from the set of shortest paths from each unassigned DC to the existing tree (P^c) until k^c datacenters are included, where $k_{min} \leq k^c \leq k_{max}$. Anycast demands for content c are routed in parallel, by connecting to the replica hosted at the closest DC. Initially, the anycast demands whose shortest paths to any DC lead to the main DC are routed as such.

The decision on which DC \hat{d} to add to the tree in each step is based on (i) the availability of storage resources needed for hosting a replica of content c , and (ii) a cost factor $\psi_{\hat{d},c}$ that is proportional to the cost associated with inter-DC multicast communication ($|p^{\hat{d},c}|$) and inversely proportional to a profit factor associated to user-driven anycast communication ($|\phi^{\hat{d},c}|$) and can be calculated as:

$$\psi_{\hat{d},c} = \frac{|p^{\hat{d},c}|}{|\phi^{\hat{d},c}|} \quad (19)$$

Here, $p^{\hat{d},c}$ is the set of links in the shortest path from \hat{d} to the existing tree $\pi^{M,c}$, while $\phi^{\hat{d},c}$ is the set of unrouted anycast demands for content c whose shortest path to any candidate DC is a path leading to \hat{d} . Note that anycast demands can be included in multiple $\phi^{\hat{d},c}$ sets for different datacenters if the lengths of the associated paths are equal. The term $|\phi^{\hat{d},c}|$ then represents the number of unrouted anycast demands that would benefit from adding a replica at candidate datacenter \hat{d} since they could be served by means of the shortest possible path to a DC.

The construction of the multicast tree for a content c starts by choosing the candidate DC \hat{d} with available storage

Algorithm 1 Lowest Cost with Adaptive Paths (MARP-H)

$\mathcal{G}(\mathcal{V}, \mathcal{E})$; \mathbf{C} : the set of contents; $\mathbf{D} \in \mathcal{V}$: the set of datacenter locations; \mathbf{W} : the set of available wavelengths; $\mathbf{k}_{\min}(\mathbf{k}_{\max})$: the minimum (maximum) number of replicas required for each content; \mathbf{S} : the maximum number of replicas a datacenter can store; \mathbf{M} : the set of manycast demands (d, c) where $d \in \mathbf{D}$ is main datacenter for content $c \in \mathbf{C}$; \mathbf{A} : the set of anycast demands (s, c) where $s \in \mathcal{V}$ and $c \in \mathbf{C}$ is the requested content;

$\mathbf{A}^c = \{\}$: the set of routed anycast demands for content $c \in \mathbf{C}$; initially, the set of demands for content c whose shortest path to any DC is a path leading to the main datacenter d ; $\bar{\mathbf{A}}^c$: the set of unrouted anycast demands for content $c \in \mathbf{C}$; $\pi^{\mathbf{A}, c} = \{\}$: the set of routes associated to anycast demands in \mathbf{A}^c ; \mathbf{S}^d : the number of contents replicas assigned to \hat{d} , initially set to zero; \mathbf{k}^c : the number of selected replica locations for content c , initially set to one; $\phi^{\hat{d}, c}$: for each $\hat{d} \in \bar{\mathbf{D}}^c$, the set of unrouted anycast demands from $\bar{\mathbf{A}}^c$ whose shortest path to \hat{d} is equal in length to the shortest path to any candidate DC; $\mathbf{D}^c = \{\mathbf{d}\}$: set of assigned DCs as replica locations for content $c \in \mathbf{C}$, initially containing only the main data center d for content c ; $\bar{\mathbf{D}}^c = \mathbf{D} \setminus \{\mathbf{d}\}$: the set of unassigned candidate DCs for content $c \in \mathbf{C}$, initially including all DCs except the main data center d for content c ; $\pi^{\mathbf{M}, c} = \{\}$: the set of links included in the manycast tree associated to assigned DCs in \mathbf{D}^c for content $c \in \mathbf{C}$; \mathbf{P}^c : the set of shortest paths $\{p^{\hat{d}, c}\}$ from each unassigned candidate DC $\hat{d} \in \bar{\mathbf{D}}^c$ to the existing tree $\pi^{\mathbf{M}, c}$; initially containing the set of shortest paths from each DC $\hat{d} \in \bar{\mathbf{D}}^c$ to the main datacenter d for content c ;

```

1: for each content  $c \in \mathbf{C}$  do
2:   while  $(k^c < k_{\max} + 1)$  do
3:     if  $(\bar{\mathbf{A}}^c \neq \{\})$  # manycast and anycast routing still need
       to be done then
4:       Select  $\hat{d} \in \bar{\mathbf{D}}^c$  with  $\min \psi_{\hat{d}, c} = |p^{\hat{d}, c}| / |\phi^{\hat{d}, c}|$  and  $S^{\hat{d}} < S$ ;
5:        $\mathbf{D}^c \leftarrow \mathbf{D}^c \cup \hat{d}$ ;  $\bar{\mathbf{D}}^c \leftarrow \bar{\mathbf{D}}^c \setminus \hat{d}$ ;
6:        $S^{\hat{d}} \leftarrow S^{\hat{d}} + 1$ ;
7:        $\mathbf{A}^c \leftarrow \mathbf{A}^c \cup \phi^{\hat{d}, c}$ ;  $\bar{\mathbf{A}}^c \leftarrow \bar{\mathbf{A}}^c \setminus \phi^{\hat{d}, c}$ ;
8:       Update  $\pi^{\mathbf{A}, c}$  by adding the shortest paths from
         anycast demands in  $\phi^{\hat{d}, c}$  to any DC in  $\mathbf{D}^c$ ;
9:       Update  $\phi^{\hat{d}, c}$  for each  $\tilde{d} \in \bar{\mathbf{D}}^c$ ;
10:      else # anycast done but manycast still in progress
11:        if  $(k^c < k_{\min} + 1)$  then
12:          Select  $\hat{d} \in \bar{\mathbf{D}}^c$  with  $\min \psi_{\hat{d}, c} = |p^{\hat{d}, c}|$  and  $S^{\hat{d}} < S$ ;
13:           $S^{\hat{d}} \leftarrow S^{\hat{d}} + 1$ ;
14:           $\mathbf{D}^c \leftarrow \mathbf{D}^c \cup \hat{d}$ ;  $\bar{\mathbf{D}}^c \leftarrow \bar{\mathbf{D}}^c \setminus \hat{d}$ ;  $k^c \leftarrow |D^c|$ ;
15:          else #  $k_{\min}$  locations selected
16:            Break
17:          end if
18:        end if
19:         $\pi^{\mathbf{M}, c} \leftarrow \pi^{\mathbf{M}, c} \cup p^{\hat{d}, c}$ ;  $\mathbf{P}^c \leftarrow \mathbf{P}^c \setminus p^{\hat{d}, c}$ ;
20:        Update the remaining paths in  $\mathbf{P}^c$  with the shortest
         paths from the new manycast tree  $\pi^{\mathbf{M}, c}$  to each unassigned
         DC in  $\bar{\mathbf{D}}^c$ ;
21:      end while
22:      if  $(\bar{\mathbf{A}}^c \neq \{\})$  # manycast (replica placement and tree) is
         done but there are still anycast demands which need to be
         routed then
23:        For remaining anycast demands  $(s, c)$  in  $\bar{\mathbf{A}}^c$ , find their
         shortest paths to any DC in  $\mathbf{D}^c$  and add to  $\pi^{\mathbf{A}, c}$ ;
24:      end if
25:      Assign the first available  $w \in \mathbf{W}$  to the constructed
         manycast tree and then to the anycast demands in order
         of their inclusion in  $\pi^{\mathbf{A}, c}$ ;
26:    end for

```

Figure 1. Pseudocode of the MARP-H algorithm.

resources and the lowest value of $\psi_{\hat{d}, c}$ (Line 4). The selected datacenter \hat{d} is added to set \mathbf{D}^c (Line 5) and the available storage capacity of \hat{d} is updated (Line 6). If more than one DC node has available storage resources and the lowest value of $\psi_{\hat{d}, c}$, the node with the lower index is selected. The anycast demands $\phi^{\hat{d}, c}$ associated with \hat{d} are then removed from the unrouted demand set $\bar{\mathbf{A}}^c$ (Line 7) and their shortest paths leading to \hat{d} are added to anycast path set $\pi^{\mathbf{A}, c}$ (Line 8). The sets $\phi^{\hat{d}, c}$ corresponding to the remaining candidate DCs are also updated by removing the newly routed anycast demands (Line 9). The links belonging to path $p^{\hat{d}, c}$ are added to the manycast tree $\pi^{\mathbf{M}, c}$ (Line 19) and the paths in \mathbf{P}^c from the remaining datacenter candidates to the existing tree are updated (Line 20).

The algorithm repeats this procedure until k^c ($k_{\min} \leq k^c \leq k_{\max}$) datacenters are included in each manycast tree, thus defining the replica placement and the associated light-trees and lightpaths. If the anycast set $\bar{\mathbf{A}}^c$ is empty before k_{\min} datacenters are included in the tree, then $\psi_{\hat{d}, c}$ is adjusted to equal $|p^{\hat{d}, c}|$, i.e., only the cost associated to inter-DC traffic defines the datacenter selection (Line 11). Conversely, if there are anycast demands left unserved after k_{\max} datacenters hosting content c are added in the tree (i.e., $\bar{\mathbf{A}}^c \neq \{\emptyset\}$), these demands are connected to the closest DC hosting a replica (Line 22). Finally, the first-fit wavelength assignment technique is used to allocate wavelengths, first to the established manycast tree and then to the anycast connections in the order in which they were added to the anycast path set ($\pi^{\mathbf{A}, c}$) for each content subsequently (Line 25).

B. An Illustrative Example for MARP-H

The operating principle of MARP-H is explained on an illustrative example shown in Fig. 2 for a single content on a European network with 16 nodes and 22 unidirectional links. The assumed main datacenter node (node 1), 5 candidate DC nodes for replica placement (nodes 3, 5, 8, 11, and 15) and 6 end-user demand source nodes for the content (nodes 6, 7, 9, 10, 12, and 13) are highlighted in the figure. In the example, we consider that content c needs to be replicated $k^c = 3$ times, i.e., $k_{\min} = k_{\max} = 3$. The initial values of $|\phi^{\hat{d}, c}|$, $|p^{\hat{d}, c}|$, and $\psi_{\hat{d}, c}$ for each candidate DC are shown in Fig. 2(a). Node 5 has the lowest cost factor (i.e., $\psi_{5, c} = 0.66$) and is, thus, added to the set of selected DCs hosting replicas in the first iteration, while the links along the shortest path from node 5 to node 1 are added to the manycast tree (shown in Fig. 2(b)). The anycast demands initiating at nodes 6, 10, and 13 are then routed to the DC at node 5, after which the cost factors for the remaining candidate DC nodes and the shortest paths from the tree to each unassigned DC node, i.e., the values of $|\phi^{\hat{d}, c}|$, $|p^{\hat{d}, c}|$, and $\psi_{\hat{d}, c}$ are updated (Fig. 2(b)). In the next iteration, nodes 3 and 8 have the lowest cost ($\psi_{3, c} = \psi_{8, c} = 1$), and node 3 is selected for hosting replica c due to its lower index. The expanded tree, the updated shortest paths, and the revised cost factors for the remaining candidate datacenters are shown in Fig. 2(c). Finally, node 8 is chosen as the third location to host a content replica while anycast demands from nodes 9 and 12 are routed towards DC at node 8. First-Fit wavelength assignment is then applied, starting with the manycast tree. The resulting routing and wavelength assignment scheme




Figure 2. An illustrative example for MARP-H for a single content c using the 16-node European topology.

for both traffic types is shown in Fig. 2(d). The procedure is done analogously for multiple contents.

C. Complexity Analysis of MARP-H

To analyze the computational complexity of MARP-H the worst-case scenario is considered, where the subroutine for selecting k_{\max} nodes for replicas of C contents, described in Lines 4-6 of Algorithm 1, has complexity of $O[k_{\max}|C||D|]$. The complexity of removing the anycast demands associated with a selected DC from the set of unrouted demands (Line 7) is $O[k_{\max}|C||A|]$. Similarly, modifying the set of unrouted anycast demands associated with the remaining candidate DCs (Lines 8-9) has complexity of $O[k_{\max}|C||A||D|]$. The complexity of updating the links included in the multicast tree (Line 19) is $O[k_{\max}|C||\mathcal{E}|]$, while the complexity of updating the shortest paths from the new multicast tree to unassigned DCs is $O[k_{\max}|C||D||\mathcal{V}|\log|\mathcal{V}|]$ (Line 20). Lastly, the complexity of first-fit wavelength assignment approach (Line 25) is $O[|C||W||\mathcal{E}|+|C||A||W||\mathcal{E}|]$.

The number of anycast demands ($|A|$), and the set of data-center locations ($|D|$) can be upper bounded by $O[|W||\mathcal{E}|]$, and $O[|\mathcal{V}|]$, respectively. Similarly, if we assume an asymptotic number of physical links of $|\mathcal{E}| \approx |\mathcal{V}|^2$ (i.e., a fully connected physical topology) as well as $|C| \ll |\mathcal{V}|$, $\log|\mathcal{V}| \ll |W|$, and $k_{\max} \ll |\mathcal{V}|$, the overall algorithm complexity can be expressed as $O[|W||\mathcal{V}|^3]$.

V. NUMERICAL RESULTS

In this section, the performance of the proposed approaches is evaluated and compared to alternative strategies




Figure 3. 14-Node NSF topology with 21 bidirectional links.




Figure 4. 30-Node European topology with 48 bidirectional links.

for manycast and anycast RWA for replica placement. To assess the effectiveness of the optimal MARP strategy, the proposed ILP formulation, MARP-ILP, which concurrently considers the inter-DC and user-to-DC traffic, as well as replica placement, is compared to two 2-step ILP-based

Table I
SIMULATION PARAMETERS AND THEIR VALUES

Parameter	Value		
	14-Node	16-Node	30-Node
Number of contents ($ C $)	10	10	20
Number of DC locations ($ D $)	6	6	8
Min(Max) number of needed replicas ($k_{\min}(k_{\max})$)	2(4)	2(4)	3(5), 5(7)
Available wavelengths (W)	24	24	128
Max number of replicas hosted in DC (S)	8	8	20

strategies which minimize the network resource usage for each type of traffic independently. The first strategy, denoted as Anycast-First ILP (AF-ILP), first finds an optimal RWA and replica placement for the anycast demands, (i.e., solving the same type of problem investigated in, e.g., [26]), and then uses this replica placement as input for finding the RWA for the inter-DC traffic. Note, this second step implies solving the multicast RWA problem (not manycast) since the destination nodes are known. The second strategy, denoted as Multicast-First ILP (MF-ILP), first solves the manycast RWA problem optimally for inter-DC traffic and then uses this replica placement to solve the anycast RWA problem for the user-driven demands. All ILP formulations are solved using the commercially available solver CPLEX on machines with Intel AMD Opteron 6172 CPUs running at 2.1 GHz. The ILP-based approaches were run for problem instances on the 14-node NSF topology (Fig. 3) and the 16-node European topology from the illustrative example (Fig. 2).

To assess the performance of the heuristic approach, MARP-H was compared to MARP-ILP for smaller problem instances, and to two benchmarking heuristics for larger problems on the 30-node European topology shown in Fig. 4. The heuristic benchmark strategies are modifications of MARP-H that adopt a 2-step approach analogous to AF-ILP and MF-ILP. The first strategy, denoted as Anycast First Heuristic (AF-H), starts by solving the anycast RWA problem and replica placement by adjusting the MARP-H DC candidate cost factor $\psi_{\hat{a},c}$ to only consider the benefit associated to anycast communication (i.e., $|\phi^{\hat{a},c}|$). Given the resulting replica placement, a multicast tree is then built in the second step by merging the paths between the main DC and the selected k^c replica hosting nodes with minimum $|p^{\hat{a},c}|$ values. The second strategy, denoted as Multicast First Heuristic (MF-H), first solves the manycast problem by adjusting the cost factor $\psi_{\hat{a},c}$ to consider only the cost of the manycast demands, i.e., $|p^{\hat{a},c}|$. The user-driven anycast demands are then connected to the closest (in terms of link count) content replica location using shortest path routing. Both AF-H and MF-H apply first-fit wavelength assignment for the established manycast light-trees and anycast lightpaths using the same policy as MARP-H. All heuristics were implemented in C, using the same machines as the ILPs.

The values of the input parameters used in the simulations are given in Table I. For each data point and parameter setting of the 14-node NSF topology and the 16-node European topology, a set of 10 test instances was generated randomly with a uniform distribution of end-user demands and candidate DC locations, solved by the ILP-based approaches and the heuristic algorithms. The same generation procedure was applied to the 30-node European network, where results of the heuristic algorithms were collected until a confidence interval of 4% or less was reached with a confidence level of 95%. The results aim to give insight into the amount of

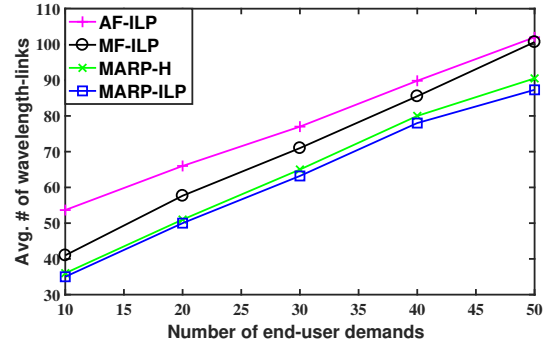


Figure 5. Total resource consumption vs. the number of end-user demands for the NSF network.

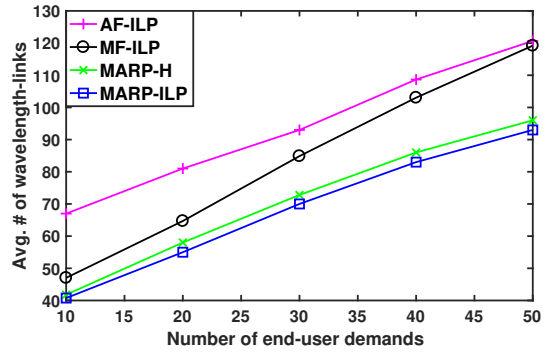


Figure 6. Total resource consumption vs. the number of end-user demands for the 16-node European network.

network capacity used by each traffic type and the possible savings that can be achieved by considering them jointly in addition to replica placement.

A. 14-Node NSF and 16-Node European Topologies

This section presents the results obtained by solving the ILP formulations and running the proposed heuristic algorithm MARP-H. Figs. 5 and 6 show the average total network resource consumption obtained by each strategy as a function of the number of end-user demands for the NSF and European topologies, respectively. When comparing the 2-step AF/MF ILPs, for a smaller number of end-user demands, the synchronization DC-to-DC traffic dominates and the manycast-first approach leads to better resource usage. As the number of user demands grows, the anycast traffic becomes prevalent and the anycast-first strategy closes the gap to the MF-ILP solutions. However, the proposed integrated MARP-ILP approach outperforms both the AF-ILP and MF-ILP strategies for all traffic loads. Specifically, MARP-ILP obtains an average resource usage reduction of 21% and 12% compared to AF-ILP and MF-ILP, respectively, for the NSF topology (Fig. 5). For the 16-node European topology (Fig. 6), these respective savings are 28% and 18%. Furthermore, the proposed heuristic, MARP-H, also outperforms the 2-step AF/MF ILP strategies and is able to obtain sub-optimal MARP solutions with an average optimality gap of less than or equal to 5%. Note that in the highest-loaded test cases, the average link utilization for both MARP-ILP and MARP-H was 28% over all links and 66% for the links adjacent to the DC nodes, indicating a significantly congested scenario.

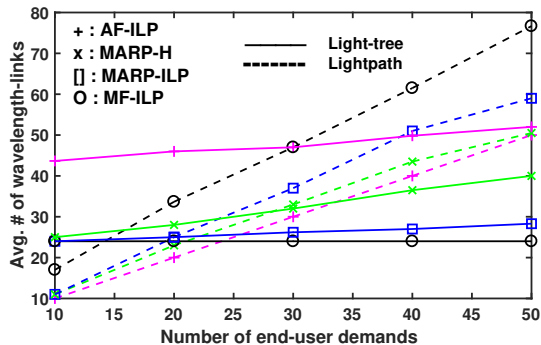


Figure 7. Resource consumption for light-trees and lightpaths vs. the number of end-user demands for the NSF network.




Figure 8. Resource consumption for light-trees and lightpaths vs. the number of end-user demands for the 16-node European network.

The performance of the MARP strategies is further analyzed by breaking down their resource usage to two distinct traffic types. Figs. 7 and 8 show the average resources allocated to the established light-trees (for synchronization/update traffic) and lightpaths (for the user-driven traffic) for the NSF and European topology, respectively. As can be expected, the AF/MF approaches that consider only one traffic type obtain the best results for that particular type, but perform poorly when it comes to the other type. Namely, AF-ILP uses the least amount of resources to establish the anycast lightpaths but yields significantly higher resource consumption for the manycast light-trees as it tends to choose relatively high values of k^c . MF-ILP, on the other hand, performs best in terms of the manycast light-trees resource usage by selecting the lowest values in the range for k^c but obtains highest resource usage for the anycast lightpaths. Due to the fact that the integrated MARP strategies (both optimal and heuristic) perform replica placement adaptively by simultaneously considering both traffic types, they can find a beneficial tradeoff between clustering the replicas to reduce the traffic resource usage for replica updates and spreading them out to diminish the user-driven resource usage. The results also indicate that the strategies establish slightly larger light-trees as the number of user-driven demands increases, in an effort to reduce the link hops for provisioned lightpaths and minimize the overall resource consumption.

The performance difference between the MARP strategies is further illustrated by Fig. 9, showing the average number of replicas per content as a function of end-user demands for the European topology. A similar pattern for the number of

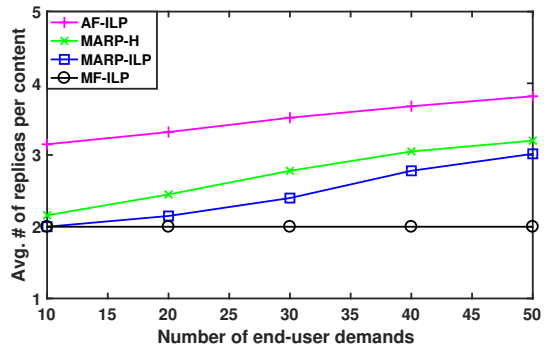


Figure 9. Average number of replicas per content vs. the number of end-user demands for the 16-node European network.

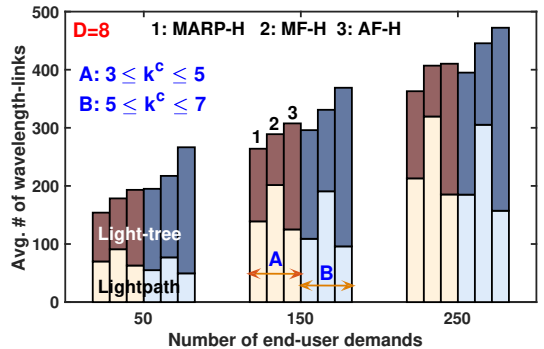


Figure 10. Resource consumption for lightpaths and light-trees vs. the number of end-user demands for the 30-node European network and $|C|=20$.

replicas selected by different MARP strategies is observed for the NSF topology and is, thus, omitted for the sake of brevity. As expected, the MF-ILP always chooses the lowest values for k^c to reduce the amount of resources used for building the manycast light-trees. On the contrary, the AF-ILP achieves the goal of storing the content closer to end-users in order to utilize the least amount of resources for the anycast lightpaths by deploying a high number of replicas per content. Between these two extremes, the integrated MARP strategies obtain a balanced value of k^c to minimize the overall resource usage generated by the two types of traffic.

Regarding execution times of the proposed algorithms, MARP-H ran under 12ms for all test cases for both topologies, while MARP-ILP required 740s to solve the largest problem instance.

B. 30-Node European Topology

This section presents the results of the MARP-H, AF-H and MF-H algorithms for the 30-node European topology under a varying number of datacenter locations and required replicas, as outlined in Table I.

Fig. 10 shows the average amount of resources consumed by the established lightpaths and light-trees in the solutions obtained for $D = 8$ DC locations, $C = 20$ contents, $\{3 \leq k^c \leq 5\}$ and $\{5 \leq k^c \leq 7\}$ required replicas per content, and a varying number of end-user demands. The results show that the proposed MARP-H strategy outperforms both MF-H and AF-H in all cases, for all values of k^c and user-demands, indicating that concurrent optimization

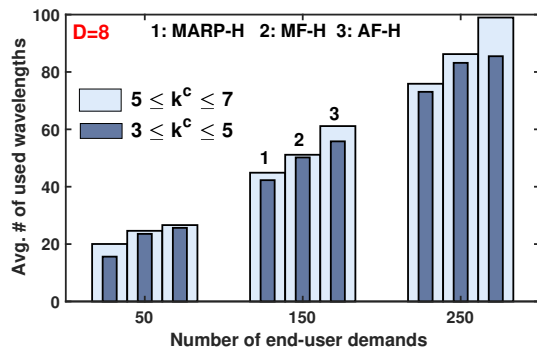


Figure 11. Avg. number of wavelengths used to establish the light-trees and lightpaths vs. the number of end-user demands for the 30-node European network and $|C|=20$.

of the two traffic types jointly with replica placement leads to enhanced solutions. On average, MARP-H saves 15% and 11% of network resources compared to AF-H and MF-H, respectively, for $3 \leq k^c \leq 5$. This saving rises to 20% and 12%, respectively, for $5 \leq k^c \leq 7$. Between AF-H and MF-H, MF-H exhibits better performance than AF-H, especially for the higher range value of k^c . The tendency of AF-H towards deploying a higher number of replicas per content leads to the creation of larger light-trees, which raises the resource usage of the DC-to-DC traffic and outweighs the benefits of having less resources consumed by the user-to-DC traffic.

Fig. 11 shows the average number of wavelengths used for the light-trees and lightpaths set up by the different strategies. MARP-H finds solutions using fewer wavelengths than the benchmarking schemes in all cases, on average by 25%. AF-H performs the worst, which can be explained by the fact it establishes larger light-trees, for which wavelength assignment under wavelength continuity constraint is more constricting than for lightpaths.

Finally, MARP-H is a very scalable alternative to solving the ILP formulation. While running slightly longer than MF-H and AF-H which ran for under 15ms, MARP-H ran for under 20ms in all the cases tested.

VI. CONCLUSION

This paper presents an integrated approach for replica placement and service provisioning in WDM-based inter-DC networks. The proposed approach applies the anycast and multicast routing paradigms to concurrently find a replica placement and the associated routing and wavelength assignment solutions for user-driven demands, as well as for updating/synchronization traffic, with the aim of minimizing the total network resource consumption. The problem is formulated as an ILP and solved optimally for smaller instances. To overcome the scalability limitation of the ILP-based optimal solution, a scalable and efficient heuristic strategy is also presented. Numerical results indicate that the proposed strategies substantially reduce network resource consumption when compared to alternative schemes which consider the two traffic types independently. Future work will include extending the model and algorithms to consider survivable and elastic optical networks.

VII. ACKNOWLEDGMENT

This work was jointly supported by the Swedish Research Council (VR) framework grant No. 2014-6230, Spanish grant

TEC2014-53071-C3-1-P (“ONOFRE”), Celtic-Plus SENDATE project, and H2020-ICT-2014 project 5GEx (Grant Agreement no. 671636).

REFERENCES

- [1] M. Habib, M. Tornatore, M. De Leenheer, F. Dikbiyik, and B. Mukherjee, “Design of disaster-resilient optical datacenter networks,” *Journal of Lightwave Technology*, vol. 30, no. 16, pp. 2563–2573, Aug. 2012.
- [2] Cisco, “Global Cloud Index: Forecast and Methodology,” 2015–2020. [Online]. Available: <https://www.cisco.com/c/dam/en/us/solutions/collateral/service-provider/global-cloud-index-gci/white-paper-c11-738085.pdf>
- [3] Mahimkar A., Chiu A., Doverspike R., Feuer M., Magill P., “Bandwidth on demand for inter-data center communication,” in *Proc. of the 10th ACM Workshop on Hot Topics in Networks*, 2011, pp. 24:1–24:6.
- [4] A. Rappaport and D. Raz, “Update aware replica placement,” in *Proc. of CNSM*, Oct. 2013, pp. 92–99.
- [5] N. Charbonneau and V. Vokkarane, “Routing and wavelength assignment of static anycast demands over all-optical wavelength-routed WDM networks,” *Journal of Optical Commun. and Networking*, vol. 2, no. 7, pp. 442–455, Jul. 2010.
- [6] A. Muhammad, N. Skorin-Kapov, and L. Wosinska, “Anycast and anycast routing for replica placement in datacenter networks,” in *Proc. of ECOC*, Sep. 2015.
- [7] B. G. Bathula, V. M. Vokkarane, C. P. Lai, and K. Bergman, “Load-aware anycast routing in IP-over-WDM networks,” in *Proc. of ICC*, Jun. 2011.
- [8] B. G. Bathula, J. M. Plante, and V. M. Vokkarane, “Crosstalk-aware anycast routing and wavelength assignment in optical WDM networks,” in *Proc. of ANTS*, Dec 2010, pp. 94–96.
- [9] X. Tan, S. Yu, and H. Zhang, “Anycast routing and wavelength assignment problem in wavelength conversion WDM networks,” *Optik - International Journal for Light and Electron Optics*, vol. 124, no. 24, pp. 6983–6986, 2013.
- [10] J. M. Plante, A. Gadkar, and V. M. Vokkarane, “Anycast overlay in split-incapable networks for supporting bandwidth-intensive applications,” *IEEE/ACM Transactions on Networking*, vol. 24, no. 1, pp. 342–354, Feb. 2016.
- [11] X. Tan, S. Yu, H. Zhang, and Y. Zhang, “An efficient algorithm for optimal anycast routing problem over WDM networks,” *Optik - International Journal for Light and Electron Optics*, vol. 124, no. 19, pp. 3866–3871, 2013.
- [12] R. Goscienc and K. Walkowiak, “Modeling and optimization of data center location and routing and spectrum allocation in survivable elastic optical networks,” *Optical Switching and Networking*, vol. 23, Part 2, pp. 129–143, 2017.
- [13] R. Lin, M. Zukerman, G. Shen, and W.-D. Zhong, “Design of light-tree based optical inter-datacenter networks,” *Journal of Optical Commun. and Networking*, vol. 5, no. 12, pp. 1443–1455, Dec. 2013.
- [14] B. Jaumard, A. Shaikh, and C. Develder, “Selecting the best locations for data centers in resilient optical grid/cloud dimensioning,” in *Proc. of ICTON*, Jul. 2012.
- [15] L. Zhang and Z. Zhu, “Dynamic anycast in inter-datacenter networks over elastic optical infrastructure,” in *Proc. of ICNC*, Feb. 2014, pp. 491–495.
- [16] Y. Chen and A. Jaekel, “Energy aware anycast routing in optical networks for cloud computing applications,” in *Proc. of CloudNet*, Oct 2015, pp. 285–290.
- [17] S. Borst, V. Gupta, and A. Walid, “Distributed caching algorithms for content distribution networks,” in *Proc. of INFOCOM*, Mar. 2010.
- [18] U. Tos, R. Mokadem, A. Hameurlain, T. Ayav, and S. Bora, “Dynamic replication strategies in data grid systems: a survey,” *The Journal of Supercomputing*, vol. 71, no. 11, pp. 4116–4140, 2015.
- [19] Y. Mansouri, S. Azad, and A. Chamkori, “Minimizing cost of K-Replica in hierarchical data grid environment,” in *Proc. of AINA*, May 2014, pp. 1073–1080.
- [20] G. Aupy, A. Benoit, M. Journault, and Y. Robert, “Power-aware replica placement in tree networks with multiple servers per client,” *Sustainable Computing: Informatics and Systems*, vol. 5, pp. 41–53, 2015.

- [21] X. Huang and Y. Peng, "A novel replica placement strategy for data center network," in *Proc. of ICITMS*, 2013, pp. 599–609.
- [22] A. Lawey, T. E. El-Gorashi, and J. Elmirghani, "Distributed energy efficient clouds over core optical networks," *Journal of Lightwave Technology*, vol. 32, no. 7, pp. 1261–1281, Apr. 2014.
- [23] I. Goiri, K. Le, J. Guitart, J. Torres, and R. Bianchini, "Intelligent placement of datacenters for internet services," in *Proc. of ICDCS*, Jun. 2011, pp. 131–142.
- [24] Y. Wu, M. Tornatore, S. Thota, and B. Mukherjee, "Renewable-energy-aware data center placement in optical cloud networks," in *Proc. of OFC*, Mar. 2015.
- [25] D. Davis, J. Plante, and V. Vokkarane, "Critical resource multi-cast protection in data center networks," in *Proc. of ICC*, Jun. 2015, pp. 5721–5727.
- [26] J. Perello, K. Walkowiak, M. Klinkowski, S. Spadaro, and D. Careglio, "Joint content placement and lightpath routing and spectrum assignment in CDNs over elastic optical network scenarios," *Computer Communications*, vol. 77, pp. 72–84, 2016.
- [27] J. Xiao, H. Wen, B. Wu, X. Jiang, P.-H. Ho, and L. Zhang, "Joint design on DCN placement and survivable cloud service provision over all-optical mesh networks," *IEEE Transactions on Commun.*, vol. 62, no. 1, pp. 235–245, Jan. 2014.
- [28] S. S. Savas, F. Dikbiyik, M. F. Habib, M. Tornatore, and B. Mukherjee, "Disaster-aware service provisioning with many-casting in cloud networks," *Photonic Network Communications*, vol. 28, no. 2, pp. 123–134, Sep. 2014.
- [29] Y. Chen, S. Jain, V. K. Adhikari, Z. L. Zhang, and K. Xu, "A first look at inter-data center traffic characteristics via yahoo! datasets," in *Proceedings of IEEE INFOCOM*, Apr. 2011, pp. 1620–1628.
- [30] R. Ravi, R. Sundaram, M. V. Marathe, D. J. Rosenkrantz, and S. S. Ravi, "Spanning trees short or small," in *Proc. of Annual ACM SIAM Symposium on Discrete Algorithms*, 1994, pp. 546–555.
- [31] D. Din, "Anycast routing and wavelength assignment problem on WDM network," *IEICE Transactions on Communications*, vol. E88B, no. 10, pp. 3941–3951, 2005.