Brokered Orchestration for End-to-End Service Provisioning across Heterogeneous Multi-Operator (Multi-AS) Optical Networks

Alberto Castro, Luis Velasco, Lluis Gifre, Cen Chen, Jie Yin, Zuqing Zhu, Roberto Proietti, and S. J. Ben Yoo, *Fellow, IEEE*

Abstract—This paper proposes a new networking paradigm introducing the broker-plane above the management planes of Autonomous Systems (ASes). The brokers communicate with the manager of each AS to assist coordinate end-to-end resource management and path provisioning across the multi-AS networks involving multiple operators. The broker plane updates the virtual network topology, manages the resource information of inter-AS links and aggregated (abstracted) intra-AS links, and computes end-to-end routing, modulation formats, and spectrum assignment (RMSA). Notwithstanding, due to the different dynamicity of each AS, the probability of finding a multi-AS transparent path fulfilling the spectrum continuity and contiguity constraints might be low. To improve the grade of the inter-domain connectivity service, spectrum converters can be installed in inter-domain nodes or per-AS defragmentation can be performed with a global view. In this paper, we introduce a mechanism where each AS can advertise its internal capabilities, e.g., spectrum conversion, their ability to implement spectrum defragmentation or any other network feature. The Multi-AS RMSA with Defragmentation Capability problem is presented and mathematically modeled. A heuristic algorithm is designed to solve it. The mechanism's workflow is comprehensively tested using simulations. Results show connection blocking reduction as high as 26%, clearly validating the benefits of the proposed mechanism. Finally, its experimental assessment was conducted on a distributed multi-continental testbed.

Index Terms—Multi-Domain Optical Networks, Network Planning, Network Brokers.

I. INTRODUCTION

SINGLE operators' transport networks are usually created as multi-domain networks a result of deploying nodes from different vendors and/or different technologies. In such scenarios, the topology of the different domains is fully visible from outside each domain and therefore it is possible to compute end-to-end paths.

Several approaches can be considered to automate end-to-end provisioning in single-operator multi-domain networks, where a Path Computation Element (PCE), possibly with other functional blocks to create an Application-Based Network Operations (ABNO) architecture [1], is in charge of

computing paths in each of the domains [2]. Each PCE (named as child PCE) has full visibility of the topology and resources of the underlying domain. A parent PCE can be connected to every child PCE to compute the sequence of domains [3], and it can compute paths for the end-to-end connection, or it can delegate intra-domain paths computation to child PCEs [4].

In contrast, as a result of privacy policies, in multi-operator multi-domain networks only an abstraction of the topologies is visible from outside the domain, which prevents from computing paths traversing more than one domain. In such scenarios, the standardized backward path computation procedure [5] allows computing an end-to-end path through a chain of PCEs, where each PCE computes a sub-path for the underlying domain. Starting from the destination PCE, the end-to-end path computation is performed towards the source PCE. Since this procedure results in sub-optimal path computation, the backward recursive PCE-based computation procedure [6] was standardized to compute optimal constrained paths. Both procedures assume that a parent PCE has previously determined the sequence of the domains. In this case, the parent PCE gets an abstracted topology of each domain that usually consists in obtaining the connectivity among interconnection nodes inside every domain.

Recent works in [7] and [8] proposed using *market-driven brokers* on top of child PCEs or Software Defined Networking (SDN) controllers in charge of each <u>Autonomous Systems (AS). Note that the interactions between the controllers and the broker are based on mutual agreements and negotiations (e.g., service level agreement) especially because each heterogeneous AS will have different requirements and agreements. In addition,</u> that scheme provides autonomy to the domains while improving scalability.

From the technological perspective, today's transport networks are mostly based on Dense Wavelength Division Multiplexing (DWDM) optical technology to exploit the huge capacity that optical connections (named as *lightpaths*) can convey. These DWDM transport networks might migrate towards Elastic Optical Networks (EONs), enabled by the flexgrid technology [9]. EONs bring flexibility as they support different spectrum allocation, as well as longer distances by using new modulation formats, e.g., Quadrature Phase Shift Keying (QPSK), which facilitates extending the transport network towards the edges [10].

In EONs, when a connection request arrives, the routing,

modulation format and spectrum allocation (RMSA) problem needs to be solved [11], [12]. One of the main causes of connection blocking in EONs is the fragmentation of the optical spectrum [13]. In the absence of optical converters, the authors proposed a re-optimization algorithm to defragment the optical spectrum that it is reactively triggered after a connection request cannot be served; the algorithm reallocates already established connections so as to make enough room for the incoming connection request. However, re-allocating a connection might entail traffic disruption that can be avoided if the standardized *make-before-break* technique is applied. In such way, reallocation procedures would not cause traffic disruption, but they require using additional network resources. To mitigate traffic disruption, authors in [14] proposed and experimentally evaluated a novel technique to perform hitless spectrum defragmentation of multiple channels without causing errors on connections lying between the initial and final spectrum locations by using a fast wavelength auto-tracking scheme. In parallel, the authors in [15] proposed the so named *push-pull* technique for hitless spectrum defragmentation; the technique performs defragmentation by moving lightpaths only to contiguous and free spectrum frequencies along the same route of the original path. Experimental validation of hitless spectrum defragmentation was reported in [16], [17] for single domain networks.

Defragmentation can also be applied in the contexts of multi-domain single operator scenarios, where the visibility of the resources within the domains is available in a centralized element. In this regard, the authors in [18] proposed a spectrum defragmentation algorithm and used an SDN controller; in addition, an inter-domain protocol to coordinate controllers in different domains was proposed.

Because of the lack of coordination among ASes, in multi-AS scenarios, transparent optical connections are rarely used. Instead, signals traversing two ASes are converted to the electrical domain and back again to the optical domain, which can be relaxed and applied only in the case of a transparent end-to-end path cannot be found.

This paper extends our previous work in [19] and study the benefits derived from applying per-domain defragmentation in multi-operator multi-domain optical networks compared to the use of optical converters. Since spectrum defragmentation is a use case of *in-operation network planning* [20], a planning tool to solve optimization problems related to network re-optimization can be deployed in each domain [21]. Furthermore, complex computations might be needed to compute end-to-end paths in multi-operator and multi-domain scenarios, and hence, a multi-domain planning tool is also proposed to that end.

The remainder of this paper is organized as follows. Section II illustrates the problem of end-to-end lightpath provisioning in multi-operator multi-AS optical networks and introduces our proposed brokered-based control architecture. Additionally, workflows for domain advertisement and lightpath provisioning are proposed. Section III is devoted to the MultiAS RMSA with Defragmentation Capability

problem, which it is first formally stated and modeled as an Integer Liner Program (ILP). Given the nature of the problem, an algorithm is proposed to solve it to optimality. Section IV presents the results obtained from exhaustive simulations, and the experimental assessment validates the feasibility of the proposed brokered architecture. Finally, Section V concludes the paper.

II. MULTI-OPERATOR MULTI-DOMAIN OPTICAL NETWORKS

This section presents our proposed brokered-based multi-operator network architecture and the workflows for *domain advertisement* and *provisioning*, including MultiAS *path computation* and *set-up*.

A. Brokered-based Multi-operator Network

Let us assume a multi-operator multi-AS flexgrid optical network, where each AS is controlled by an SDN/OF controller or an ABNO-based architecture. On top of the ASes, a broker coordinates end-to-end multi-AS provisioning and includes a planning tool for complex computations (Fig. 1).

Since both, the broker and the planning tool will be requested to perform complex computations, each AS is assumed to advertise inter-domain nodes and links (while updating the available spectrum for each inter-domain link to follow updates) independently from path computation requests (Fig. 2a). In addition, each AS might agree to expose further features, named as *capabilities*, which can be supported by specific hardware (e.g., spectrum converters) or by optimization algorithms (e.g., spectrum defragmentation).

When a computation is requested, the broker collects intra-AS abstracted connectivity and spectrum availability (Fig. 2b). Observe that, each AS advertises an abstracted intra-AS link information to the broker that depends on both, internal AS policies and the specific agreement with the broker. Details of the AS intra-domain topology remains concealed from the rest ASes and the broker.

A solution might entail applying a capability in an AS. For example, using spectrum converters or performing defragmentation to release a set of slices so that an end-to-end lightpath can be established. In the example in Fig. 2c, the AS200 SDN controller is requested to use conversion from slot 1 to slot 2, while in Fig. 2d the controller applies defragmentation in its domain to release frequency slots 1 or 6. Note that by doing so, end-to-end lightpaths can be established.

The next section defines the notation used along the paper.

B. Notation

The broker has a global view of the virtualized network topology, including information on inter-AS links and abstracted intra-AS link status gathered from each AS.

To represent the underlying data plane, we use a graph $G(N_O, E_O)$, where N_O is the set of optical nodes and E_O is the set of optical links connecting two nodes. Graph G is structured as a set A of ASes. Every AS $a \in A$ consists of three differentiated subsets of nodes:
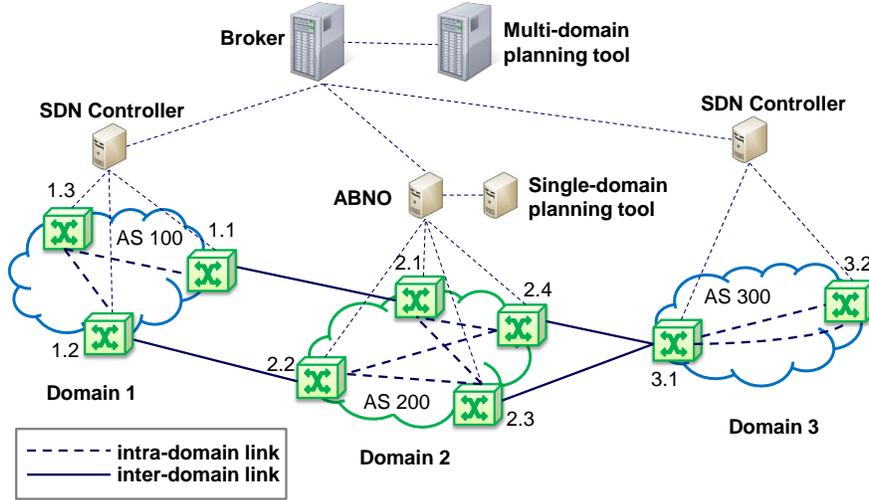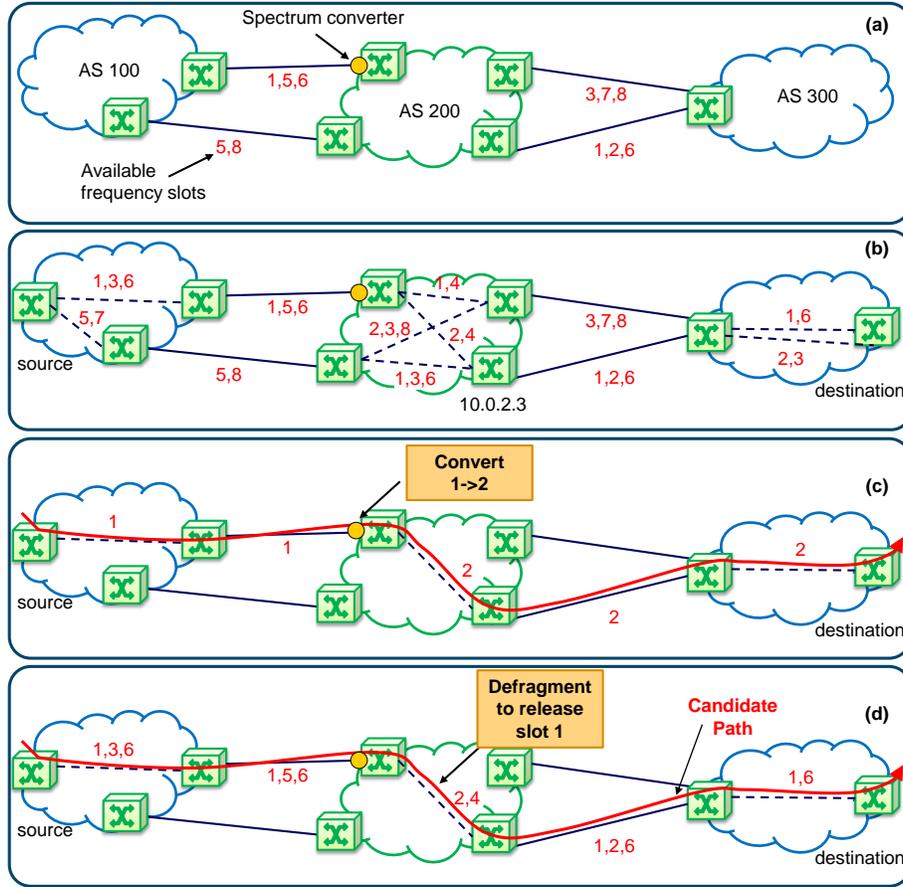
Fig. 1. Multi-AS network architecture.

Fig. 2. An example of path computation. a) Initial AS advertisement and inter-AS link updating. b) Abstracted intra-AS and end-nodes advertisement. c) Path computation using the conversion capability. d) Path computation using the defragmentation capabilities.

$N_e(a)$ subset of edge nodes belonging to AS a, end-points of demands.

$N_t(a)$ subset of AS a transit (internal) nodes.

$N_i(a)$ subset of AS a border nodes.

Then, $N_{(\cdot)} = \bigcup_{a \in A} N_{(\cdot)}(a)$, and $N_O = N_e \cup N_t \cup N_i$, with $N_e \cap N_i = \emptyset$. For instance, $N_i(\text{AS100}) = \{1.1, 1.2\}$ and $N_e(\text{AS100}) = \{1.3\}$ in Fig. 1.

Regarding the links, two subsets are considered:

$E_n(a)$ subset of abstracted intra-AS a links. Each $e \in E_n(a)$ abstracts connectivity between either a node in $N_e(a)$ and another node in $N_i(a)$ belonging to the request's end ASes, or between two nodes in $N_i(a)$ belonging to intermediate ASes.

$E_i(a)$ subset of inter-AS links starting in a node in $N_i(a)$ and ending in a node in $N_i(a')$. $E_i = \bigcup_{a \in A} E_i(a)$

Finally, let S be the set of available frequency slices in each optical link e. Each e is then represented by a tuple $<u_e, v_e, S_e,$

c_e>, where u_e, $v_e \in N_e \cup N_i$ are the end nodes, S_e is the subset of available frequency slices, and c_e is the metric (cost).

The next section defines the workflows to implement the proposed brokered orchestration.

C. Workflows definition

Fig. 3 illustrates the proposed workflows. First, the Domain Advertisement workflow is initiated when the broker first connects to the ASes controllers (message 1 in Fig. 3). The broker collects inter-AS information, along with the AS's capabilities. The broker updates the inter-AS topology to keep databases aligned (message 3).

Second, the provisioning workflow, which includes the *Path Computation* and the *Path Set-up* phases, is shown. The path computation phase is triggered by the arrival of a new inter- AS path computation request to an SDN controller. The SDN controller forwards the request to the broker (message 5), which collects intra-AS connectivity (messages 6-7) from every AS and applies its local RMSA algorithm [11], [12] to find a transparent end-to-end lightpath.

Table I presents the algorithm for Intra-AS abstract connectivity collection and path computation. Starting from the input graph G(N_i, E_i), abstract inter-AS links E_n is collected. Abstract connectivity from source and destination ASes is collected first (lines 2-3 in Table I) and then, from every other AS (lines 4-5). Each SDN controller computes one or more paths and available slots between every node in the first input set and every other node in the second input set; every path might be associated with a different cost. A new graph G' is created by adding source and destination nodes to the set of nodes and the collected inter-AS links to the set of links (line 6). A transparent lightpath <p, c_p> is eventually computed (line 7) on G', being p the ordered subset of links in its route and c_p, the selected frequency slot.

If the broker succeeds, the new lightpath is set-up.

TABLE I. INTRA-AS CONNECTIVITY COLLECTION AND PATH COMPUTATION ALGORITHM

IN: G(N_i, E_i), d
OUT: p, c_p
1:     **if** $d.src.a = d.dst.a$ **then return** Ø
2:     $E_n(d.src.a) \leftarrow$ get_connectivity($d.src.a$, {$d.src$}, $N_i(d.src.a)$)
3:     $E_n(d.dst.a) \leftarrow$ get_connectivity($d.dst.a$, $N_i(d.dst.a)$, {$d.dst$})
4:     **for each** a in $A \backslash \{d.src.a, d.dst.a\}$ **do**
5:         $E_n(a) \leftarrow$ get_connectivity(d, $N_i(a)$, $N_i(a)$)
6:     G' $\leftarrow$ G($N_i \cup \{d.src, d.dst\}$, $E_i \cup E_n$)
7:     **return** RMSA (G', d)

Otherwise, the broker makes a path computation request to the planning tool, adding the just collected topology information to the request message (message 8). Upon the reception, the planning tool runs the algorithm in Table II.

If the planning tool finds a feasible solution that entails using capabilities in some AS, it responds `NO-PATH` but proposes testing one or more capabilities in the ASes (message 9). In such case, the broker requests testing whether the selected capabilities are still available, e.g., whether a frequency slot can be released between a pair of nodes by applying defragmentation (message 10). If the capabilities are successfully tested, the broker sends a new path computation request to the planning tool allowing the possibility of using the tested capabilities during the computation (message 12). Eventually, the planning tool responds with the multi-AS path to be set-up and the list of capabilities to be used (message 13).

Finally, the broker, following the solution proposed by the planning tool, instructs the SDN controllers to signal the intra-AS path and configure the border routers (messages 14-15). Once all the SDN controllers finish their local set-up, the broker informs the initiating SDN controller that the inter-AS path is available (message 16).
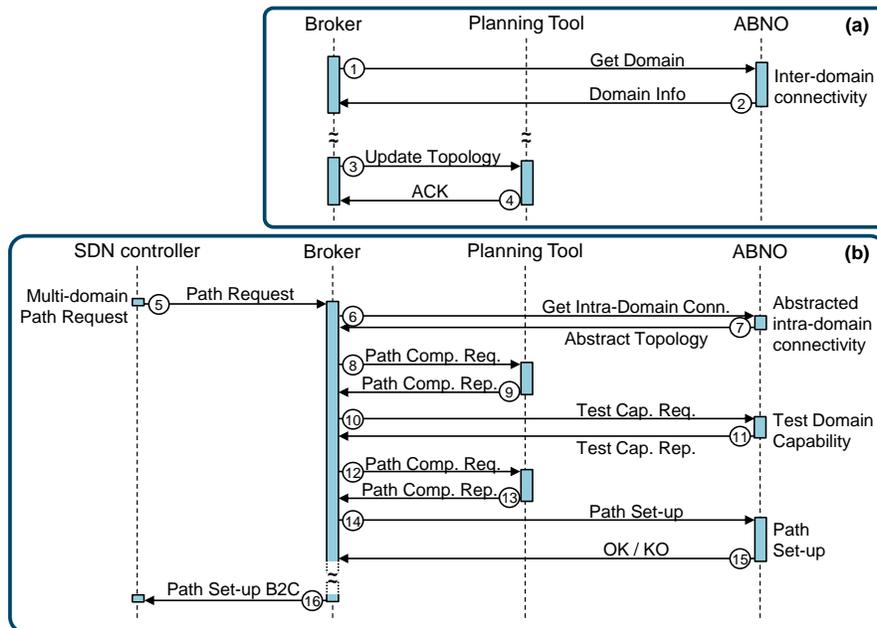

Fig. 3. Proposed workflows. a) Domain advertisement and b) Path computation and set-up.

TABLE II. CAPABILITY-BASED MULTIAS_RMSA ALGORITHM

IN: G(N, E), d	**OUT:** p, c_p, AC

1:	$Q \leftarrow \emptyset$
2:	$C \leftarrow$ getSlots(d.bw)
3:	$P = \{p\} \leftarrow$ kSP(G, d.src, d.dst)
4:	**for each** p **in** P **do**
5:	bestSA ← <-1, ∞, Ø>
6:	**for each** c **in** C **do**
7:	cost ← 0; AC← Ø
8:	**for each** e **in** p **do**
9:	**if** free(e.c) **then**
10:	cost ← cost + e.cost
11:	**continue**
12:	**if** $e \in E_i$ OR NOT defrag $\in$ A(e).capabilities **then**
13:	cost ← ∞
14:	**break**
15:	cost ← cost + e.cost + e.defrag.cost
16:	AC ← AC ∪ A(e)
17:	**if** cost < bestSA.cost **then** bestSA ← <c, cost, AC>
18:	**if** bestSA.cost < ∞ **then** Q ← Q ∪ {<p, bestSA>}
19:	**if** Q = Ø **then return** Ø
20:	sort(Q, <q.cost, ASC>, <|q.p|, ASC>)
21:	**return** first(Q)

The next section models the MultiAS RMSA with Defragmentation Capability problem and proposes an algorithm to solve it to optimality.

III. MULTIAS RMSA WITH DEFRAGMENTATION CAPABILITY

This section presents solving methods for the in-operation planning algorithm that selects a transparent end-to-end multi-AS lightpath that will be feasible after applying announced capabilities in a subset of ASes. The problem can be formally stated as follows.

Given:
- a connected graph G(N, E). Each link e with a given cost.
- the optical spectrum and the used modulation format,
- a demand d, specifying the source and destination nodes and the required bitrate

Output: the route and spectrum allocation for d and the ASes where the defragmentation capability needs to be applied.

Objective: minimize the cost of serving d.

The graph G(N, E) is built with the abstracted inter-AS links E_n received from the broker.

In the following, we present an ILP formulation to solve the above problem, based on the formulations in [11]. Note that since the topology is given, we can pre-compute a set P of distinct paths for the demand. Moreover, for the sake of clarity, one single modulation format is considered, e.g., QPSK. Thus, the set of frequency slots can be computed. The following sets and parameters have been defined.

Topology:
- N   Set of optical nodes, index n.
- E   Set of fiber links, index e. Each link with a transmission cost c_e and a defragmentation cost k_e. $k_e = \infty$ if $e \in E_i$ or the AS represented by e does not support the defragmentation capability.

Paths and Spectrum:
- P   Set of pre-computed paths, index p. Each path with a cost $c_p = \Sigma_e \, \delta_{pe} \cdot c_e$.
- S   Set of spectrum slices, index s.
- C   Set of pre-computed slots for the demand.
- δ_{pe} Equal to 1 if path p uses link e.
- δ_{cs} Equal to 1 if slot c uses slice s.
- α_{es} Equal to 1 if slice s in link e is free.

The decision variables are:
- x_{pc} Binary, equal to 1 if demand is routed through path p and slot c; 0 otherwise.
- y_e Binary, equal to 1 if the defragmentation capability needs to be applied to the AS represented by the abstracted link e; 0 otherwise.

The ILP formulation is as follows:

$$\min \quad \sum_{p \in P}\sum_{c \in C} c_p \cdot x_{pc} + \sum_{e \in E} k_e \cdot y_e \tag{1}$$

subject to:

$$\sum_{p \in P}\sum_{c \in C} x_{pc} = 1 \tag{2}$$

$$\sum_{p \in P}\sum_{c \in C} \delta_{pe} \cdot \delta_{cs} \cdot x_{pc} - y_e \le \alpha_{es}, \quad \forall e \in E, s \in S \tag{3}$$

The objective function (1) minimizes the cost of serving a demand regarding the path cost and AS(es) to defragment. Constraint (2) ensures that a lightpath, i.e., a path and a frequency slot, is selected for the demand. Constraint (3) stores whether defragmentation needs to be applied in link e to release slot c. Note that if variable y_e is activated for a link where defragmentation cannot be applied, the problem is infeasible and the demand is blocked.

Because we are dealing with one single demand, the problem can be solved to the optimality using the algorithm presented in Table II. The set of frequency slots is computed (line 2 in Table II) while the kSP algorithm finds a number k of shortest routes between source and destination (line 3). For each computed route, the best frequency slot is found (lines 4-18). Every frequency slot is evaluated in all the links of the route (lines 8-16); if the slot is free in a link, the route cost is increased with the cost of that link. In the opposite, if the link is inter-AS or the AS has not announced the defrag capability, the route is discarded; otherwise, the cost of the defragmentation capability in the AS is added to the route cost. The set AC is updated with the ASes where the defragmentation capability needs to be applied. Feasible routes are stored in the set Q (line 18). If the set Q is empty, there is no solution and the demand will be blocked (line 19). Otherwise, the set Q is sorted by route cost first and then by the length of the route, both in ascending order, (line 20) and the best route is eventually returned (line 21).

The next section reports the performance evaluation and the experimental tests carried out to validate the feasibility of the proposed brokered orchestration.

IV. RESULTS

In this section, we first evaluate the performance of the proposed per-domain spectrum defragmentation scheme. Illustrative simulation results are presented from using the MultiAS RMSA with Defragmentation Capability algorithm over a realistic multi-operator multi-domain network topology. Next, the feasibility of the proposed architecture is experimentally assessed in a real environment.

A. Performance Evaluation

Performance evaluation was carried out on a multi-domain scenario with three optical network topologies (Fig. 4); the 22-node 35-link British Telecom (BT), the 21-node 35-link Spanish Telefonica (TEL), and the 22-node 38-link Telecom Italia (TI) topologies. In addition, two inter-domain links connecting BT and TEL (BT-1 - TEL-9 and BT-11 - TEL-1), and other two connecting TEL and TI (TEL-20 - TI-7, and TEL 19 - TI-14) have been considered. Regarding capabilities, the TEL domain implements the *defragmentation* and the spectrum conversion ones.

For evaluation purposes, we developed an ad-hoc event-driven simulator in OMNET++ [22]. A dynamic network environment was simulated where incoming connection requests for both intra-domain and inter-domain traffic arrive at the system following a Poisson process and are sequentially served without prior knowledge of future incoming connection requests. To compute the RMSA of the intra-domain lightpaths, we used the algorithm described in [12]. The holding time of the connection requests is exponentially distributed with a mean value equal to 2 hours. Source/destination pairs are randomly chosen with equal probability (uniform distribution) among all nodes. It is worth noting that every node in the three domains can be source or destination of *intra-domain* connections in contrast to *inter-domain* connections, where only nodes in BT and TI networks can be source or destination. Different values of the offered network load are created by changing the inter-arrival rate while keeping the mean holding time constant. We assume that no retrial is performed; if a request cannot be served, it is immediately blocked. Regarding the optical spectrum, we assumed a total width of 4 THz with a spectrum granularity of 6.25 GHz.
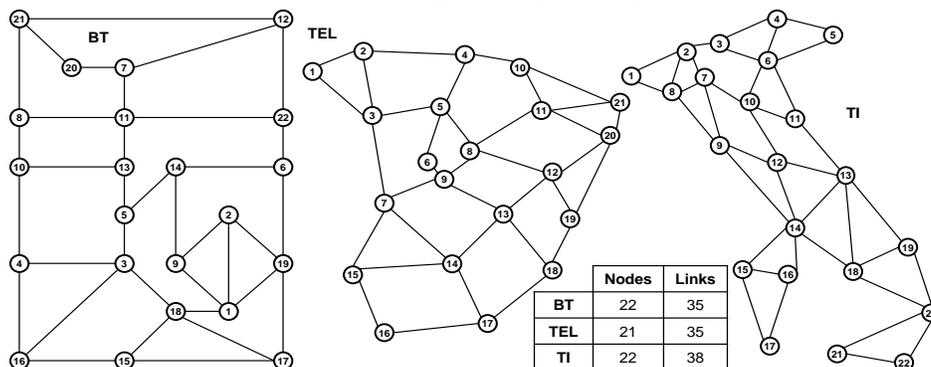
In our simulations, the bitrate of every connection request was set to 100 Gb/s. To convert bitrate into spectrum width, we use the correspondence in [23]. Finally, each point in the results is the average of 10 independent runs with 15,000 connection requests each, where the first 5,000 were used for warming-up purposes.

Fig. 5 presents the blocking probability of every domain against the offered load when only intra-domain connections are requested. We observe that the load that every topology can carry for a given target blocking probability is different. In consequence, in the rest of this section, we identify the different loads by the resulting blocking probability and study the impact of inter-domain connections when every domain is loaded, so its intra-domain blocking probability is 1%, 0.5%, and 0.1% (dotted lines in Fig. 5).

Graphs in Fig. 6 plot inter-domain blocking probability against the inter-domain offered load when no capabilities are used (labeled as *Transparent*) and when the defragmentation capability is used (labeled as *Defragmentation*) for the three selected intra-domain loads. The load has been normalized with respect to 150 Erlangs. As observed, the inter-domain blocking probability highly depends on the inter-domain load, and even for moderate loads, the inter-domain blocking probability reaches unacceptable values. This high blocking is a consequence of the improbability of finding continuous spectrum in the three domains simultaneously. Although applying the defragmentation capability brings benefits as high as 300% under the 1% intra-domain-blocking load (Fig. 6a) and 260% under the 0.5% one (Fig. 6b), many connection requests cannot still be served. It is not until the intra-domain load is reduced to 0.1% intra-domain blocking probability (Fig. 6c), that inter-domain blocking reduces to tolerable levels. There, the benefits of using the per-domain defragmentation capability allow conveying 26% more traffic that when no capability is utilized for a target inter-domain blocking of 0.1%. We also show that the impact of inter-domain connections on the intra-domain blocking probability is negligible as a result of the low inter-domain load when compared to the intra-domain one.

Fig. 4. Network topologies considered: the 22-node British Telecom (BT), the 21-node Telefonica (TEL), and the 22-node Telecom Italia (TI).

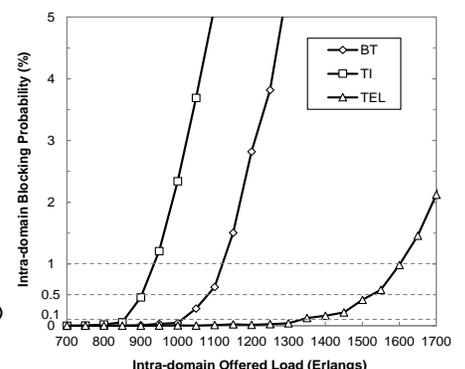|	Nodes	Links
BT	22	35
TEL	21	35
TI	22	38

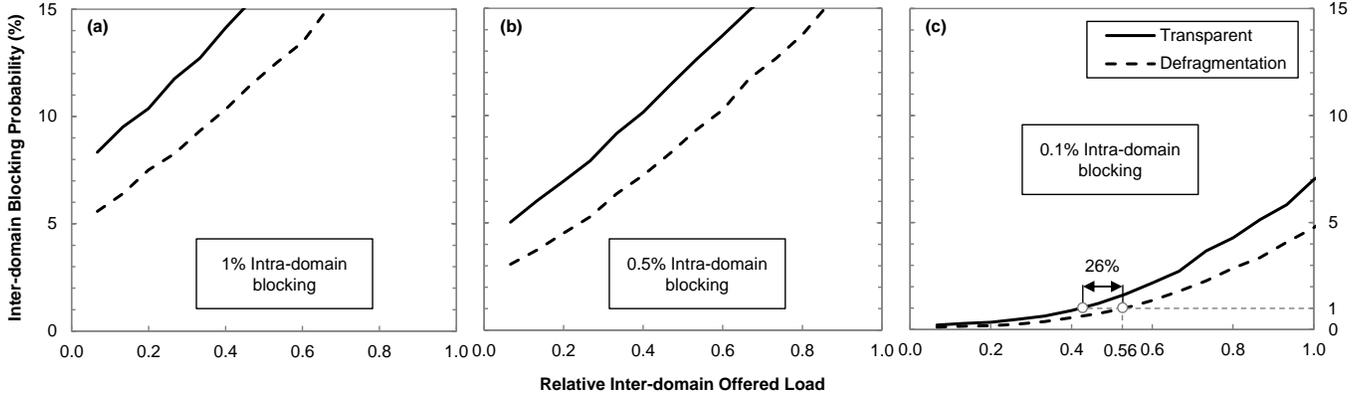Fig. 5. Intra-domain blocking probability against offered load.

Fig. 6. Inter-domain blocking probability against relative intra-domain offered load, when domains are loaded so the resulting intra-domain blocking probability is (a) 1%, (b) 0.5%, and (c) 0.1%.
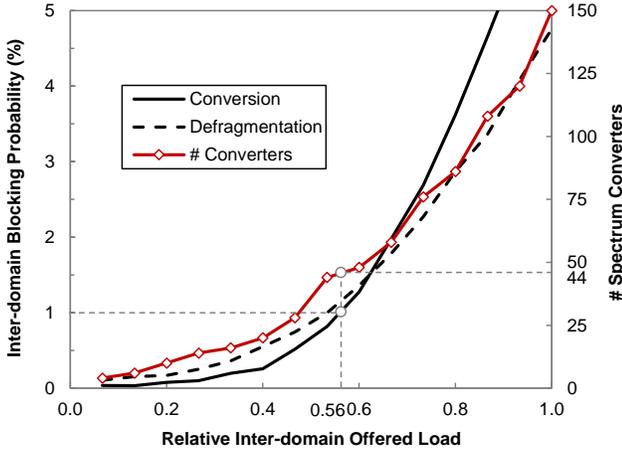

Fig. 7. Performance of Conversion and Defragmentation capabilities.

TABLE III. BLOCKING REDUCTION FROM APPLYING DEFRAGMENTATION

Load (%)	Transparent	Defragmentation	Reduction
1.0	16.5%	12.5%	**24%**
0.5	12.6%	9.3%	**26%**
0.1	1.6%	1%	**39%**

Table III presents the blocking probability reduction obtained for the inter-domain normalized load equals 0.56 for the three considered inter-domain loads.

Fig. 7 shows a comparison of inter-domain blocking as a function of the normalized inter-domain load for an intra-domain load equals 0.1% when Conversion and Defragmentation capabilities are used; the total number of spectrum converters are plotted as well. As observed, using the conversion capability provides some benefits in terms of additional load that can be conveyed. Note that those benefits are virtually zero for the load that results in a 1% inter-domain blocking probability, where the amount of spectrum converters that are needed is as high as 44, which discourages using the conversion capability.

B. Experimental Assessment

The experimental validation was carried out on a distributed field trial set-up spanning three continents connecting premises in UC Davis (Davis, California), USTC (Hefei, China), and UPC (Barcelona, Spain) (Fig. 8).

The broker, the OF controllers, and agents were developed in Python and run in a computer cluster under Linux.

The UPC's SYNERGY test-bed includes the multi-domain PLAnning Tool for Optical Networks (PLATON) and the ABNO, developed in C++ for Linux. The implementation of the ABNO architecture consists of: *i*) the ABNO controller as the entrance point to the optical transport network for provisioning and advanced network coordination; *ii*) an active stateful Path Computation Element to serve path computation requests able to modify established connections, e.g., for defragmentation purposes; *iii*) an in-operation planning tool, a different instance of PLATON deployed as a dedicated back-end PCE. The PCE Communication Protocol (PCEP) is used to carry path computation requests and responses.

Regarding the management plane, to enable the broker to orchestrate the experiment, we developed an HTTP REST API at the broker, which is implemented by the SDN controllers and PLATON. For each API function, a specific XML has been devised; these XML messages act as input/output parameters for the API functions.

Fig. 9 shows the exchanged messages from the broker viewpoint. For the sake of clarity, message numbering used in the workflows have also been included. The domain advertisement workflow starts when the broker connects to the SDN controllers and populates its topology. Every time a new topology is obtained, a copy is sent to PLATON, so as to maintain broker and PLATON databases synchronized (messages 1-4). Fig. 10 shows details of some selected messages. Specifically, message 2 contains the set of nodes and inter-AS links in AS 200; spectrum availability is encoded as a hexadecimal number. Interestingly, AS advertises capability "40" (defragmentation).

Message 5 in Fig. 10 gives details of the message received by the broker in the event of a path computation request from an SDN controller. The broker starts collecting abstracted intra-AS connectivity from every AS; Fig. 10 shows the details of message 7 received from AS 100. After running its local RMSA algorithm, the broker cannot find a feasible transparent end-to-end lightpath, so it decides to send an in-operation planning request to PLATON. PLATON first
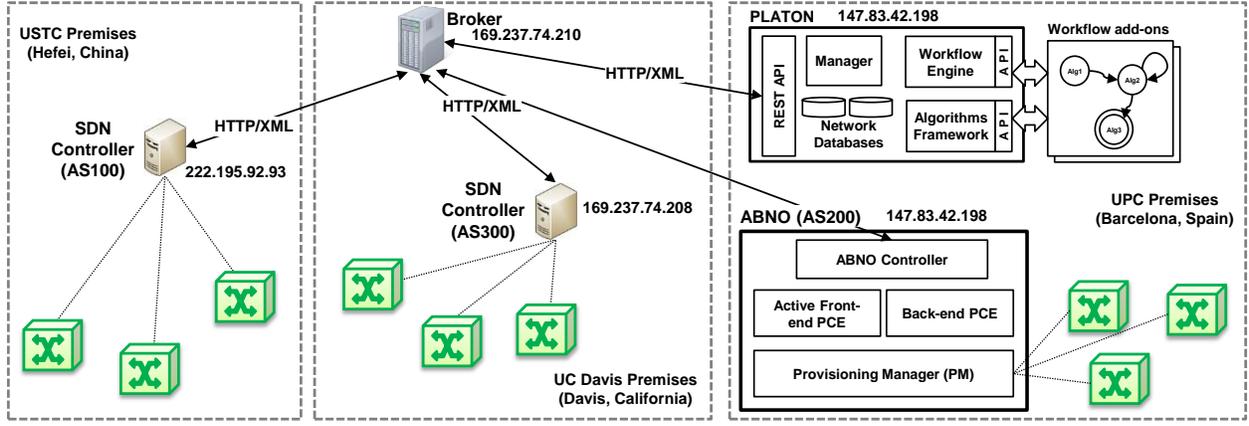
Fig. 8: Experimental set-up

Fig. 9: Messages Exchange at the broker

Fig. 10: Detail of selected XML messages

updates its database with the topology information contained in the incoming request message, and then it runs the algorithm in Table II.

As a result of the spectrum availability configured in our set up, no solution is found. Consequently, a NO-PATH reply is sent to the broker (see message 9 in Fig. 10). Within the reply, PLATON suggests that if defragmentation is applied to abstracted link .2.1-.2.3 in the AS 200 to release slice 1, a solution can be found.

Next, the broker accepts PLATON suggestion and requests testing the defragmentation capability to the ABNO controller in AS 200 (see message 10 in Fig. 10). Let us assume that the result of the test is positive, i.e. the slice can be released by defragmentation, so the ABNO controller replies OK. Right after, the broker resends a path computation request to PLATON, but this time announcing that the defragmentation capability can be applied on link .2.1-.2.3 in the AS 200 to release slice 1. Now PLATON can find a solution that is sent to the broker (see message 13 in Fig. 10). The XML contains the routing and spectrum allocation, and the capability to be applied.

Next, the broker creates the set of configurations to be forwarded to the corresponding SDN controllers/ABNO.

When every controller confirms that the configuration has been set-up, the broker informs the requester SDN controller that the multi-AS path has been signaled.

V. CONCLUSIONS

A broker on top of opaquely-managed domains was extended to perform per-domain spectrum defragmentation when no feasible transparent end-to-end lightpath can be found for a multi-domain connectivity request.

During their startup, every domain announces their border nodes, their inter-domain connectivity, and their capabilities, e.g., spectrum defragmentation and conversion, to the broker. When a multi-domain connection request is received, the broker tries to compute a transparent end-to-end lightpath, and if no feasible lightpath can be found, it requests to a multi-domain in-operation planning tool (PLATON) to propose a solution possibly applying some of the capabilities announced by the domains. The MultiAS RMSA with defragmentation capability problem was formally stated and an ILP formulation was proposed. An exact algorithm was devised to solve the problem to optimality.

Extensive simulations were carried out on a three domain network to evaluate the performance of the per-domain

spectrum defragmentation scheme. Results showed that an increment of the inter-domain conveyed traffic as high as 26% when the intra-domain traffic was kept at 0.1% of blocking probability. When the conversion capability was used, a significant amount of spectrum converters were needed to reach the same results.

Finally, the proposed scheme was experimentally validated on an inter-continental distributed field trial set-up. Two SDN controllers and an ABNO controller were in charge of the opaquely-managed domains. The broker on top of them, assisted by PLATON, was responsible for the domain orchestration and the spectrum defragmentation capability was enabled in the ABNO-controlled domain.

REFERENCES

[1] D. King and A. Farrel, "A PCE-based Architecture for Application-based Network Operations," IETF RFC 7491, 2015.

[2] F. Paolucci, F. Cugini, A. Giorgetti, N. Sambo, and P. Castoldi, "A Survey on the Path Computation Element (PCE) Architecture," IEEE Communications Surveys & Tutorials, vol. 15, pp. 1819-1841, 2013.

[3] D. King and A. Farrel, "The Application of the Path Computation Element Architecture to the Determination of a Sequence of Domains in MPLS and GMPLS," IETF RFC6805, 2012.

[4] O. González de Dios et al., "Multipartner Demonstration of BGP-LS-Enabled Multidomain EON Control and Instantiation with H-PCE," IEEE/OSA Journal of Optical Communications and Networking, vol. 7, pp. B153-B162, 2015.

[5] N. Bitar, R. Zhang, and K. Kumaki, "Inter-AS Requirements for the Path Computation Element Communication Protocol (PCEP)," IETF RFC5376, 2008.

[6] JP. Vasseur, R. Zhang, N. Bitar, and JL. Le Roux, "A Backward-Recursive PCE-Based Computation (BRPC) Procedure to Compute Shortest Constrained Inter-Domain Traffic Engineering Label Switched Paths," IETF RFC5441, 2009.

[7] S. J. B. Yoo, "Multi-domain Cognitive Optical Software Defined Networks with Market-Driven Brokers," in Proc. European Conference on Optical Communication (ECOC), 2014.

[8] Dan Marconett and S. J. B. Yoo, "FlowBroker: Market-Driven Multi-Domain SDN With Heterogeneous Brokers," in Proc. IEEE/OSA Optical Fiber Communication Conference (OFC), 2015.

[9] M. Ruiz, L. Velasco, A. Lord, D. Fonseca, M. Pióro, R. Wessäly, and J.P. Fernández-Palacios, "Planning Fixed to Flexgrid Gradual Migration: Drivers and Open Issues," IEEE Communications Magazine, vol. 52, pp. 70-76, 2014.

[10] L. Velasco, P. Wright, A. Lord, and G. Junyent, "Saving CAPEX by Extending Flexgrid-based Core Optical Networks towards the Edges," (Invited Paper) IEEE/OSA Journal of Optical Communications and Networking (JOCN), vol. 5, pp. A171-A183, 2013.

[11] L. Velasco, A. Castro, M. Ruiz, and G. Junyent, "Solving Routing and Spectrum Allocation Related Optimization Problems: from Off-Line to In-Operation Flexgrid Network Planning," IEEE/OSA Journal of Lightwave Technology (JLT), vol. 32, pp. 2780-2795, 2014.

[12] O. González de Dios et al., "Experimental Demonstration of Multi-vendor and Multi-domain Elastic Optical Network with data and control interoperability over a Pan-European Test-bed," IEEE/OSA Journal of Lightwave Technology, vol. 34, pp. 1610-1617, 2016.

[13] A. Castro, L. Velasco, M. Ruiz, M. Klinkowski, J. P. Fernández-Palacios, and D. Careglio, "Dynamic Routing and Spectrum (Re)Allocation in Future Flexgrid Optical Networks," Elsevier Computers Networks, vol. 56, pp. 2869-2883, 2012.

[14] C. Qin, R. Proietti, B. Guan, Y. Yin, R. Scott, R. Yu, and S. J. B. Yoo, "Demonstration of Multi-channel Hitless Defragmentation with Fast Auto-tracking Coherent RX LOs," in Proc. IEEE/OSA Optical Fiber Communication Conference (OFC), 2013.

[15] F. Cugini, F. Paolucci, G. Meloni, G. Berrettini, M. Secondini, F. Fresi, N. Sambo, L. Potí, and P. Castoldi, "Push-pull defragmentation without traffic disruption in flexible grid optical networks," IEEE/OSA Journal of Lightwave Technology (JLT), vol. 31, pp. 125–133, 2013.

[16] F. Paolucci, A. Castro, F. Fresi, M. Imran, A. Giorgetti, B. Bhownik, G. Berrettini, G. Meloni, F. Cugini, L. Velasco, L. Potí, and P. Castoldi, "Active PCE Demonstration performing Elastic Operations and Hitless Defragmentation in Flexible Grid Optical Networks," Springer Photonic Network Communications (PNET), vol. 29, pp. 57-66, 2015.

[17] Ll. Gifre, F. Paolucci, A. Aguado, R. Casellas, A. Castro, F. Cugini, P. Castoldi, L. Velasco, and V. López, "Experimental Assessment of In-Operation Spectrum Defragmentation," Springer Photonic Network Communications, vol. 27, pp. 128-140, 2014.

[18] Z. Zhu, X. Chen, C. Chen, S. Ma, M. Zhang, L. Liu, and S. J. B. Yoo, "OpenFlow-Assisted Online Defragmentation in Single-/Multi-Domain Software-Defined Elastic Optical Networks," (Invited paper) IEEE/OSA Journal of Optical Communications and Networking (JOCN), vol. 7, pp. A7-A15, 2015.

[19] A. Castro, Ll. Gifre, C. Chen, J. Yin, Z. Zhu, L. Velasco, and S. J. B. Yoo, "Experimental Demonstration of Brokered Orchestration for end-to-end Service Provisioning and Interoperability across Heterogeneous Multi-Operator (Multi-AS) Optical Networks," in Proc. ECOC, 2015.

[20] L. Velasco, D. King, O. Gerstel, R. Casellas, A. Castro, and V. López, "In-Operation Network Planning," IEEE Communications Magazine, vol. 52, pp. 52-60, 2014.

[21] Ll. Gifre, F. Paolucci, L. Velasco, A. Aguado, F. Cugini, P. Castoldi, and V. López, "First Experimental Assessment of ABNO-driven In-Operation Flexgrid Network Re-Optimization," IEEE/OSA Journal of Lightwave Technology (JLT), vol. 33, pp. 618-624, 2015.

[22] OMNET++: http://www.omnetpp.org/

[23] Ll. Gifre, R. Martínez, R. Casellas, R. Vilalta, R. Muñoz, and L. Velasco, "Modulation Format-aware Re-Optimization in Flexgrid Optical Networks: Concept and Experimental Assessment," in Proc. European Conference on Optical Communication (ECOC), 2015.