



HAL
open science

Virtual Link Embedding in Collaborative Sliced Multi-Administrative Multi-Domain Networks

Stanislas Pedebearn, Slim Abdellatif, Pascal Berthou, Dariusz Nogalski,
Dallal Belabed

► **To cite this version:**

Stanislas Pedebearn, Slim Abdellatif, Pascal Berthou, Dariusz Nogalski, Dallal Belabed. Virtual Link Embedding in Collaborative Sliced Multi-Administrative Multi-Domain Networks. SIGAPP Symposium On Applied Computing, ACM, Apr 2024, Avila (ES), Spain. hal-04502307v2

HAL Id: hal-04502307

<https://laas.hal.science/hal-04502307v2>

Submitted on 15 Apr 2024

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

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

Virtual Link Embedding in Collaborative Sliced Multi-administrative Multi-Domain Networks

Stanislas Pedebearn
LAAS-CNRS, Université de Toulouse,
CNRS, UPS, Toulouse, France
spedebearn@laas.fr

Slim Abdellatif
LAAS-CNRS, Université de Toulouse,
CNRS, UPS, Toulouse, France
slim@laas.fr

Pascal Berthou
LAAS-CNRS, Université de Toulouse,
CNRS, UPS, Toulouse, France
berthou@laas.fr

Dariusz Nogalski
Military Communication Institute,
Zegrze, Poland
Dariusz.Nogalski@wil.waw.pl

Dallal Belabed
Airbus Defence and Space,
Issy-les-Moulineaux, France
dallal.belabed@airbus.com

ABSTRACT

When providing services in multi-administrative multi-domain networks, domains usually disclose a topology aggregation composed of a set of abstract links connecting their border nodes. Such an aggregation is motivated by the will of limiting the information exposed to other domains. This, however, leads to inefficient resource usage. In order to enable effective collaborations between domains, in this paper, we propose to enrich the topology aggregation exposed by domains by promoting the inclusion of abstracted non-border network nodes and the network slices supported by each domain. Our evaluations on real and random topologies show the significant potential gains in terms of admission ratios and resource usage that such additions bring.

KEYWORDS

Link embedding, multi-administrative multi-domain networks, topology aggregation, network slicing, QoS

1 INTRODUCTION

When embedding an end-to-end multi-domain service with some Quality of Service (QoS) guarantees, the initiating domain relies on the aggregated topologies exposed by other domains (i.e., compact portray of the network topology of each domain) to combine multiple portions and compose the end-to-end service. Classically, to restrict the topology information that is disclosed, the aggregated topology that domains expose to the others is composed of border nodes fully or partially connected to each other with abstract links (we name it below as the legacy aggregated /abstracted topology).

We argue that in some contexts, domains may need to closely collaborate to achieve a shared ultimate goal. Such goals can be: optimally spreading the overall load on the different domains in order to improve the admissibility of forthcoming end-to-end service demands, minimizing resource usage in all domains, minimizing the end-to-end service provisioning time, etc. This is for instance the case in federated military networks where ally nations connect and share

some of their network infrastructure to build a multi-domain Federated Mission Network (FMN) with the goal of leveraging on this multi-domain FMN to provide the communication services needed during the military mission [7]. In this case, the ultimate goal is to efficiently use the resources shared between nations and maximize the support of the end-to-end services needed for the mission. A similar situation exists in the civilian context since service providers usually have privileged partners; their respective domains may also work in close collaboration.

We argue that collaborative domains need to expose more abstracted topology information than the legacy topology aggregation. Indeed, this latter hides lots of details and leads either to an overestimation or underestimation of the available resources. This in turn leads to less efficient resource utilization, more service demand denials, etc.

Network domains that are adopting network slicing may support multiple domain-level slices, each providing predefined types of service that can be exposed and made available to other domains in order to compose end-to-end multi-domain services. In such cases, slices can also be disclosed as part of the aggregated topology with the characterization of the type of service provided by the slice. In this work, we focus on transport slices, i.e., slices providing connectivity with some predefined QoS between exposed nodes. One example of such a slice is a Low-latency slice, which interconnects a set of nodes exposed by the domain with a transfer delay of a couple of tenths of milliseconds. Another example is a slice providing efficient point-to-multipoint transmission services within the slice boundaries.

In this paper, we address the problem of embedding end-to-end virtual links in collaborative multi-domain networks, where network domains are allowed to expose an aggregated topology that combines border nodes and their connecting abstract links, and also, non-border abstract nodes and domain-level slices. In this perspective, we evaluate the benefit of these two latter components when compared to the legacy aggregated topology.

This paper is organized as follows. Section 2 reviews the state of the art. Then, by considering an example of a real network topology, we describe the motivations behind the proposed topology aggregation and then define it. Based on this, section 4 revisits the virtual link embedding problem in a collaborative multi-domain network context. Section 5 evaluates the performance impact of the proposed topology aggregation in comparison to the legacy one.

2 STATE OF THE ART

Virtual link and virtual network embedding have been extensively researched during the last decades [6] by implicitly assuming a single administrative domain with a thorough knowledge of the substrate network. With the advent of network slicing and VNF service chaining, some recent research is considering the context of multi-administrative multi-domain networks [10]. To the best of our knowledge, most of the proposed algorithms commit to restricting the information (domain topology and network resources) that is disclosed among domains by adopting the legacy abstraction with a mesh of links between edge nodes (peering nodes). Houidi et al. [8] expose full substrate topology but split the substrate network model into two types of information: functional (static) attributes available to virtual network (VN) providers e.g., node type, geographic location, cost, and non-functional attributes (dynamic: capacity, actual QoS) - not available to VN providers (VNP), but available only to infrastructure providers (InP) (i.e., operator of the domain). Shen et al. [9] assume that three types of resource information in each substrate domain are provided by the InP to VNP, including: the location of the edge nodes, capacity, price; inter-domain link capacity, price; intra-domain links remain under cover so as to protect the detailed topology of each InP from VNP and its competitors, but each edge (peering) node to edge node price is provided. Dietrich et al. [5], additionally to the information about links between the InP's peering nodes (the bandwidth cost), assume that each InP discloses also each offered virtual node type (resource type) along with the associated cost.

Several papers discuss the related problem of service function (network function) placement. Toumi et al. [11] follows similar to [9] information sharing scheme for multi-domain SFC placement with limited visibility. Abujoda and Papadimitriou [1] solve a problem of distributed network service embedding, where only peering links between multiple providers are exposed. Dietrich et al [4] solve multi-provider service chain embedding. Authors derive the abstract network view that obscures the Internet access points (available in substrate) and includes (i) a full mesh of DCs (data center) and peering nodes within each NFP (network function provider) (each link connecting a DC with a peering node or a pair of peering nodes corresponds to the shortest path between the respective nodes), and (ii) the peering among NFPs.

In this paper, we take the opposite direction to above works and investigate the benefit of an enriched domain abstraction.

3 TOPOLOGY AGGREGATION IN COLLABORATIVE SOFTWAREZED MULTI-DOMAIN NETWORKS

This section introduces the topology of a Polish network domain illustrated in Figure 1. It then introduces some topology contraction rules and delve into some intuitive methods of contracting a network topology with the associated trade-offs.

3.1 MOTIVATING EXAMPLE

Consider the topology of Figure 1, which describes the real topology of a Polish network domain taken from [2] with its physical nodes and links. The labels attached to the links represent their capacity in units of bandwidth; Link delays are omitted for clarity; The border nodes of the domain (edge nodes) are represented in green, i.e., nodes E1, E2, and E3; Last, the red clouds correspond to clients that can be the initiator or one of the destinations of an end-to-end virtual link demand. Various aggregated topologies may be derived based on the aggregation policy employed by the domain. Below, we elaborate on an intuitive method, which leads to different aggregated topologies that can be exposed by the domain. We discuss the advantages and flaws of each obtained topology. At first, the idea is to successively identify elementary local topologies that can be contracted to one node such that any set of demands that can be admitted by the contracted subgraph is also admissible with the non-contracted subgraph and vice versa.

Figure 2 illustrates a first rule, which shows that any set of bandwidth demands assigned on the link (n1,n2) can also be supported on the link (E3,n1) in both directions. From a bandwidth allocation perspective, nodes E3 and n1 can be merged into an abstracted node A2. The capacity of the directional links between this abstracted node and n2 (right-hand figure) captures the bandwidth constraints on all the demands that flow between nodes E3 and n2. This rule can be applied twice on the topology of Figure 1 (on E3 and the top left corner). Figure 3 illustrates the second contraction rule that applies to the group of three nodes. All bandwidth demands that meet the bandwidth constraints of the external links that connect the group of three nodes to other nodes can be supported by bandwidth allocations on internal links (connecting the three nodes). From

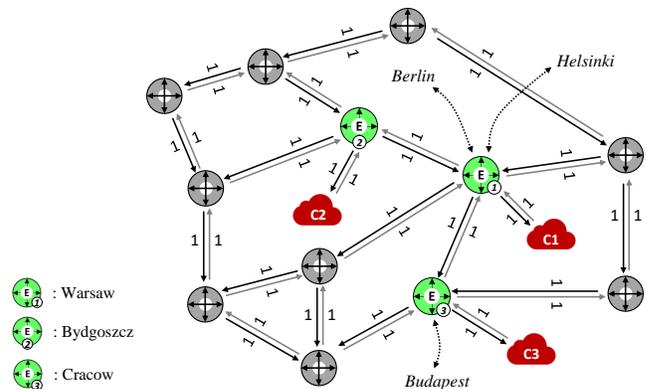


Figure 1: Polish network domain.

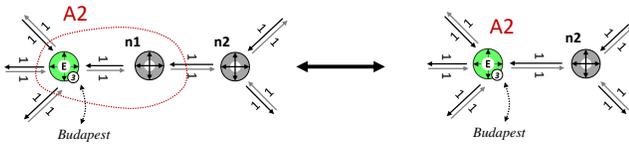


Figure 2: Illustration of contraction rule 1.

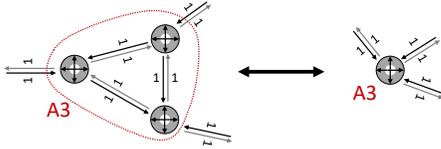


Figure 3: Illustration of contraction rule 2.

a resource allocation perspective, the three nodes can be merged into an abstract node A3. When applying rule 2 three times, the contracted topology of Figure 4 is obtained. It is worth to note that any demand that can be supported by the contracted topology is also admissible by the non-contracted initial topology of Figure 1. Also, the obtained graph includes a non-border node e.g., node A1, which is an abstracted node obtained after topology contraction. With our proposed topology aggregation, the contracted graph of Figure 4 can be disclosed by the considered domain. If the policy of the domain requires the exposure of the legacy aggregated topology with only border nodes, additional contraction of the preceding graph becomes necessary. Two options can be considered:

- Exposing an over-provisioned topology : A topology that advertises more resources than what is available. In such cases, some demands may be accepted at the multi-domain level, but when they go through the domain-level admission control for confirmation, they are rejected; An over-provisioned aggregated legacy topology can be obtained by artificially increasing the link (A1,E3) capacity in both directions to 2 units. The resulting aggregated topology that can be exposed by the considered domain is depicted in left-hand Figure 5.
- Exposing an under-provisioned topology : A topology that advertises less resources than what is available, ensuring that any demand that is admitted at the multi-domain level is admitted at the domain level. In such cases, some demands may be rejected uselessly at the multi-domain level while sufficient resources are available in the domain. An under-provisioned aggregated topology can be obtained by removing link (E1,A1) from the topology of Figure 3. The resulting aggregated topology that can be exposed by the considered domain is depicted in right-hand Figure 5.

3.2 PROPOSED TOPOLOGY AGGREGATION

The aggregated topology that a domain exposes is classically composed of a list of border nodes with a list of links connecting them. These latter are typically abstract links and correspond to one or multiple data paths (with multiple physical hops) followed by the traffic flowing from one border node to the pair node. Each has an SLA associated to it, which specifies the expected performance

in terms of maximum and available capacity, maximum transfer delay, maximum packet loss rate, etc. Links may also be labeled with cost information or any other information that can be used to enforce some policy-based decisions. The proposed aggregation paves the way for the inclusion of additional network nodes with their associated links to capture some topological constructs, and also, the network slices supported and made available by the domain. Even if they are typically not physical (i.e., abstracted), non-border nodes are considered as switching nodes and share the same characteristics as border nodes. We consider the following attributes to characterize the exposed slices:

- Entry & Exit points: It lists the exposed nodes' interfaces that are connected to the slice, i.e., from where packets will benefit from the service provided by the slice; Slice entry or exit points are characterized by:
 - Maximum flow rate, which defines the maximum packet flow rate that is allowed when respectively entering or leaving the slice at the considered point.
 - Access delay, when needed. It reflects the fact that an excess delay may be needed to capture the time required to reach the slice from the entry/exit point.
- Slice type: Specifies the type of network services that the slice is providing. Slice types are shared and agreed between all domains. We assume two different families of slice types: (1) Transport slices, which provide raw network connectivity services with a predefined SLA; (2) SFC (Service Function Chaining) based slices, which provide network services that include network capabilities defined in the SFC with a predefined SLA. Without losing generality, in the mathematical formulation below, we only consider transport slices.
- SLA: Expressed, as classically, in terms of max transfer delay, packet loss rate, etc.

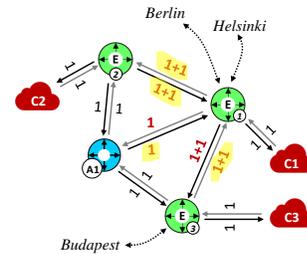


Figure 4: Aggregated topology with non-border node.

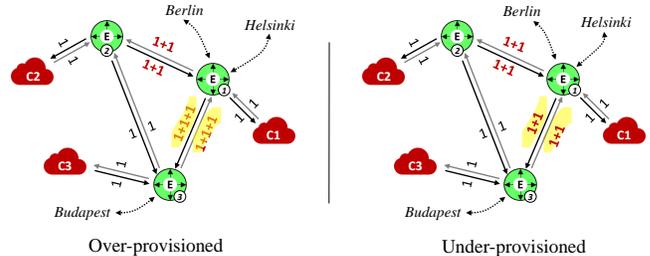


Figure 5: Aggregated topology with over/under-provisioning.

- Maximum throughput: Specifies the maximum amount of incoming traffic (whatever the entry point) that the slice can handle, i.e., the sum of the rates of all packet flows entering the slice must not exceed this capacity.
- Provisioning time: Indicates the time required to provision the slice on the infrastructure of the domain.

3.3 ABSTRACTED MULTI-DOMAIN GRAPH

Aggregated topologies are collected and gathered by each domain, then used as input to build their own view of the aggregated multi-domain topology, derive their multi-domain resource graph, which is then used as input of their resource allocation algorithm.

We adopt a classical representation of the abstracted multi-domain resource graph. Border and non-border nodes are represented as vertexes with switching resources expressed as a forwarding table size and a forwarding rate. We also model each exposed slice as a graph vertex that inherits the forwarding rate of the slice. Links connect the different types of vertexes. They are classically characterized with classical performance metrics, e.g., delay, capacity, etc. For the links that enter or leave a slice, their bandwidth is respectively set to the maximum flow rate of the corresponding entry or exit point of the slice. The max transfer delay is reflected by the incoming links of the slice. Access delays are assigned to incoming and outgoing links of the slice. For illustration, let us reconsider the network of Figure 1 and assume that the domain is also providing a low latency slice with a max throughput of 2 units of bandwidth and a max transfer delay of 50 ms. We also assume that all border nodes are connected to the slice, and that node E3 requires an extra delay of 10 ms to reach the slice. Figure 6 depicts the resulting topology aggregation with the labels assigned as explained above (for clarity, delays of edge-to-edge links are hidden).

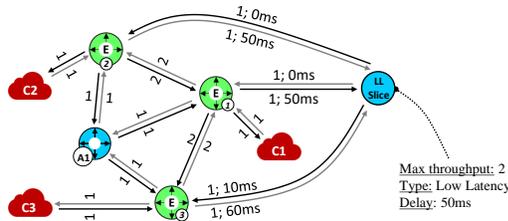


Figure 6: Abstracted topology with abstracted slice.

4 MATHEMATICAL FORMULATION

The Abstracted multi-domain resource graph, which is used for resource allocation does not exhibit any peculiarities that would refrain from applying any algorithm devised for the single-domain case. Below, we revisit the resource allocation algorithm of [3] by including the notion of abstracted non-border node and slice node. For space reasons, we partially describe the mathematical formulation derived from the multi-commodity flow problem. We notably do not present the aspects related to the allocation of switching resources. Nonetheless, they are subject to very similar constraints and optimization objectives as those related to bandwidth allocations. We also do not present the path-splitting capability of the

algorithm. Requests arrive and are treated in sequence with no information on future requests (i.e., online). Each demand gathers multiple concurrent end-to-end virtual link requests, each concerning a point-to-point or point-to-multipoint virtual link with bandwidth and delay requirements. By default, the objective of the resource allocation is to distribute fairly network traffic and use efficiently network resources and, when needed, to favor the use of slices at the expense of border node-to-border node links. Below, the multi-domain resource graph and virtual links demand models are described; Then, the variables and problem constraints are listed; Lastly, the considered objective function is defined.

4.1 MULTI-DOMAIN RESOURCE GRAPH MODEL

It is a unidirectional graph $G' = (V', A')$ where $V' = N_t \cup N_s$ is the set of disclosed network nodes (border nodes + non-border nodes) and disclosed network slices, and $A' (|A'|, A' \subseteq V' \times V')$ is the set of exposed abstract links connecting exposed network nodes to each other or to exposed slices, and, lastly, inter-domain links. For each slice $i \in N_s$, we respectively denote by FR_i^{\max} and FR_i the maximum and remaining slice throughput. Lastly, for each link $(i, j) \in A'$, we respectively denote by $\gamma_{i,j}^{\max}$ and $l_{i,j}$, the maximum capacity of the link and the maximum packet transfer delay along the link.

4.2 DEMAND MODEL

Each demand is composed of a set of D end-to-end virtual unidirectional links. Each virtual link k is characterized by:

- A source node $o_k \in N_t$, and a set of destination nodes $T_k \subseteq N_t - \{o_k\}$ (when $|T_k| = 1$, the virtual link is point-to-point; otherwise, it is point-to-multipoint).
- A bandwidth requirement of b_k , a maximum transfer delay of L_k , and a maximum packet size of p_k .

4.3 RESOURCE-RELATED ASSIGNMENT VARIABLES

The output of the resource allocation considered in this section is the set of data paths (with the bandwidth allocations at each supporting substrate link) that support each of the virtual links composing the demand. Since virtual links may be point-to-multipoint variables are related to a specific destination of a virtual link as follows:

- $\phi_{i,j}^{k,t}$: represents the bandwidth assigned to the packets of virtual link k that flow from the source node o_k to a destination node t from T_k at link $(i, j) \in A'$. Also, $\phi_{i,j}^k$ refers to the amount of bandwidth used on link (i, j) by the virtual link k regardless of the destination. It is set to the maximum of $\phi_{i,j}^{(k,t)}$ for all destinations $t \in T_k$.
- $x_{i,j}^{k,t}$: is a Boolean variable which reflects whether the flow of packets of virtual link k destined to t is supported by link $(i, j) \in A'$ ($x_{i,j}^{k,t} = 0$ if $\phi_{i,j}^{k,t} = 0$ and $x_{i,j}^{k,t} = 1$ otherwise). Variables $x_{i,j}^{k,t}$, $t \in T_k$, are used to derive the Boolean variable $x_{i,j}^k$, which indicates if some bandwidth from link $(i, j) \in A'$ is assigned to k .

4.4 PROBLEM CONSTRAINTS

The constraints on bandwidth allocations are described in equations (1) to (5). Inequalities (1) reflect the linearization of the Max operator applied to variables $\phi_{i,j}^{k,t}$ to get $\phi_{i,j}^k$.

$$\begin{aligned} \forall k \in D, \forall (i, j) \in A', \forall t \in T_k : \quad & \phi_{i,j}^{k,t} \leq \phi_{i,j}^k \\ \forall k \in D, \forall (i, j) \in A' : \quad & \phi_{i,j}^k \leq \sum_{t \in T_k} \phi_{i,j}^{k,t} \end{aligned} \quad (1)$$

Inequalities (2) ensure that the bandwidth assigned to each virtual link k at link (i, j) does not exceed the remaining bandwidth. Equation (3) represents the usual flow conservation constraints.

$$\forall (i, j) \in A' : \quad \sum_{k \in D} \phi_{i,j}^k \leq Y_{i,j} \quad (2)$$

$$\forall k \in D, \forall t \in T_k, \forall i \in V' :$$

$$\sum_{\substack{j \in V' \\ (i,j) \in A'}} \phi_{i,j}^{k,t} - \sum_{\substack{j \in V' \\ (j,i) \in A'}} \phi_{j,i}^{k,t} = \begin{cases} b_k & \text{if } i = o_k \\ -b_k & \text{if } i = t \\ 0 & \text{else} \end{cases} \quad (3)$$

Inequality (4) is a channeling constraint between integer and Boolean variables: $\phi_{i,j}^k$ and $x_{i,j}^k$. Additionally, they impose a constraint on the bandwidth assignment of virtual link k at a substrate link, restricting it to the requested bandwidth b_k .

$$\forall k \in D, \forall (i, j) \in A' : \quad x_{i,j}^k \leq \phi_{i,j}^k \quad \text{and} \quad \phi_{i,j}^k \leq b_k \cdot x_{i,j}^k \quad (4)$$

As part of the aggregated topology exposed by a domain is the maximum forwarding rate of incoming packets, which holds for network nodes and slices. Inequality (5) ensures that bandwidth allocations respect these max throughput.

$$\forall i \in V' : \quad \sum_{k \in D} \sum_{(j,i) \in A'} \phi_{j,i}^k \leq FR_i \quad (5)$$

Virtual links have end-to-end delay requirements. For point-to-multipoint links, they must be met for all of their end destinations. These requirements are considered by inequalities (6). As path-splitting is not considered in this paper, the transfer delay that a packet from virtual link k destined to t experiences is the accumulation of the delays experienced at each hop. Each hop, through either a transport node or slice node, induces a packet transmission time, bounded by $\frac{p_k}{b_k}$ and switching and propagation delays bounded by latency $l_{i,j}$ for any link $(i, j) \in A'$. Indeed, as explained in section 3.3, the slice delay is reported on incoming links.

$$\forall k \in D, \forall t \in T_k : \quad \sum_{(i,j) \in A'} x_{i,j}^k \cdot \left(\frac{p_k}{b_k} + l_{i,j} \right) \leq L_k \quad (6)$$

4.5 OBJECTIVE FUNCTION

The objective function aims at minimizing link and node resource consumption but also at distributing the consumed resources among nodes and links in order to reduce the creation of bottlenecks. To that end, referring to inequalities (7), variable φ_{\max} captures the maximum link utilization (when considering all network links) after demand D acceptance. φ_{\max} being minimized by the objective function.

$$\forall (i, j) \in A' : \quad 1 - \left(\frac{1}{Y_{i,j}^{\max}} \cdot (Y_{i,j} - \sum_{k \in D} \phi_{i,j}^k) \right) \leq \varphi^{\max} \quad (7)$$

Expression (8) defines our general and tunable objective function. It consists of three components, each weighted with a parameter that controls the impact of the component on the resolution process. In the first term, the aspect that is minimized is the average utilization of network links after allocating bandwidth to the considered demand D . In the second term, the aspect that is minimized is the maximum network link utilization. This means that the allocations devoted to the request are distributed over different links in such a way that link load disparity is reduced. When activated, the last term favors the use of slices at the expense of border node-to-border node links (with constants $c_1 < c_2$).

$$\begin{aligned} & \alpha_1 \cdot \frac{1}{|A'|} \cdot \sum_{(i,j) \in A'} \left(\frac{1}{Y_{i,j}^{\max}} \cdot (Y_{i,j}^{\max} - Y_{i,j} + \sum_{k \in D} \phi_{i,j}^k) \right) \\ & + \alpha_2 \cdot \varphi^{\max} \\ & + \alpha_3 \cdot \left(c_1 \cdot \sum_{\substack{(i,j) \in A' \\ j \in N_s}} x_{i,j}^k + c_2 \cdot \sum_{\substack{(i,j) \in A' \\ j \in N_t}} x_{i,j}^k \right) \end{aligned} \quad (8)$$

5 PERFORMANCE EVALUATION

The objective is to quantitatively evaluate under different topologies and service requests, the impact of the proposed topology aggregation on the resource allocation of multi-domain end-to-end virtual links, as described in section 4. More precisely, two different aspects are evaluated below:

- The benefits of allowing the disclosure of non-border network nodes with respect to the legacy topology aggregation;
- The benefits of disclosing domain-level slices when compared to a multi-domain aggregation without slices.

5.1 EVALUATION OF THE IMPACT OF THE PROPOSED TOPOLOGY AGGREGATION

5.1.1 Simulation set-up.

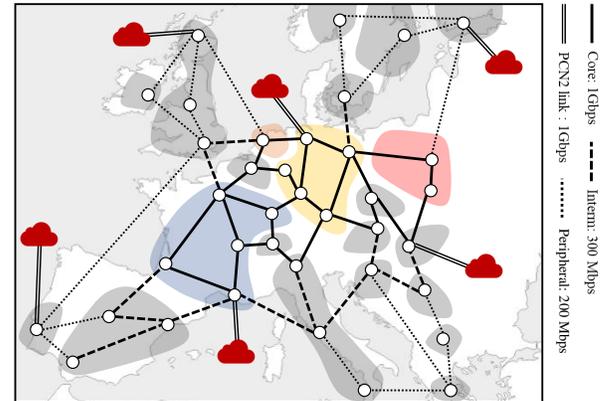


Figure 7: COST 266 based multi-domain network.

To consider a realistic scenario, we consider the Cost266 topology [2] presented in Figure 7. It is composed of 22 domains from different countries. Apart from the English domain, all other domains expose exclusively border nodes with border-to-border node links. In most cases, border nodes from the same domain are not connected in full mesh. The capacities of exposed links are presented in Figure 7. Out of the 22 domains, we vary the Polish and French aggregated topology by considering different legacy aggregations and the one promoted in this paper. The physical topologies of the Polish and French domains are respectively presented in Figure 1 and on the left-hand Figure 9 and are taken from [2]. As the information on links is not available, we assume that all links have the same capacity of 1 unit. Also, we assign 9 clients to the French domain. By connecting the different versions of the Polish and French aggregated topologies to the 20 other domains, four multi-domain topologies are hence considered for the evaluation:

- Topology “proposed”: The Cost266 topology with the proposed aggregated topologies of the Polish and French domains as shown in Figure 1 (Polish) and right-hand Figure 9 (French). For both domains, any demand that can be supported by each of the aggregated topology is admissible by the corresponding physical topology;
- Topology “over-p”: The Cost266 topology with the over provisioned legacy aggregation of the Polish and French as respectively presented in left-hand Figure 4 and Figure 8;
- Topology “under-p1”: The Cost266 topology with a first option of the under-provisioned legacy aggregation of the Polish and French domain, as respectively presented in right-hand Figure 4 and right-hand Figure 8. This latter is obtained by removing links (L,13) and (B,M) (right-hand Figure 8) before the final contraction;
- Topology “under-p2”: The Cost266 topology with a second option of the under-provisioned legacy topology aggregation of the Polish and French domain as respectively presented in right-hand Figure 4 and Figure 8. This latter is obtained by reducing the capacity of links (L,13), (L,M), (P,13) and (B,M).

Without loss of generality, service demands are composed of one single end-to-end link (which can have one or multiple destinations) with bandwidth and delay requirements. The number of service demands is varied in order to consider different loads (from low/medium loads up to very high loads (with 1000 demands)). The source and destination of the end-to-end link are randomly generated from the set of French and Polish clients.

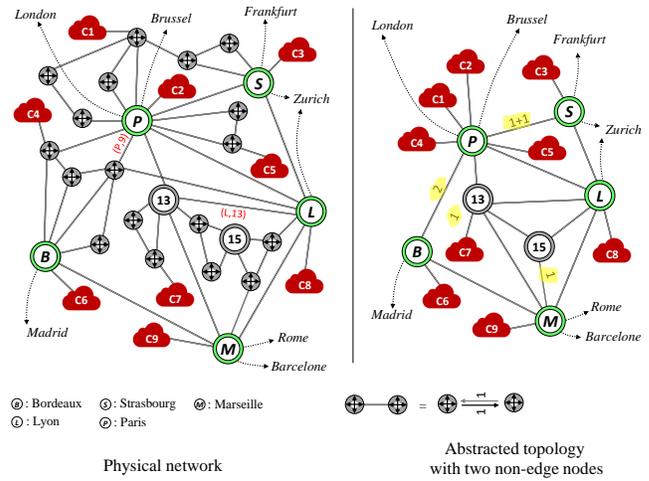


Figure 9: Physical and aggregated French network.

5.1.2 Simulation set-up.

Figure 10 describes the admission ratio as a function of time for the four topologies, as well as the corresponding number of admitted demands at the end of the simulation. As expected, the overprovisioned legacy topology aggregation exhibits the best admission ratio with, at the end of the simulation, 9 more accepted demands when compared to the aggregation proposed in this paper. It is worth to note that these 9 demands are in reality not feasible and even if they are considered by the resource allocation of the multi-domain level (i.e., using the overprovisioned legacy aggregated), the resource allocation running at the domain level will reject them. As a consequence, a new admission will need to be triggered to compute a new allocation. In other words, the overprovisioned topology aggregation induces computing/network overhead as well as extra delays in addressing a service demand. These could be avoided with the proposed aggregation. Figure 10 also shows that when using an under-provisioned legacy aggregation, at the end of the simulation, more than 10% of the demands are rejected while in fact, they are feasible (as shown with the “proposed” abstraction). Hence, resources are wasted. Indeed, while available, they are not assigned to arriving demands because of an unprecise abstraction.

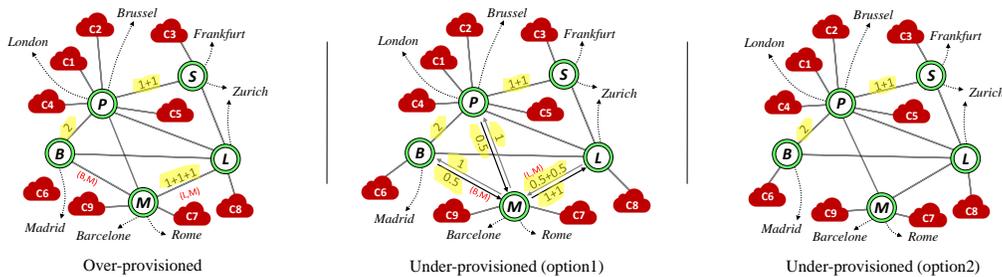


Figure 8: Over-provisioned and under-provisioned (option 1 and 2) legacy abstraction of the French domain.

Our simulations show that these observations are also valid when load balancing is activated. As the topology aggregation of the French domain is trickier to derive, below, we assume that all the demands are originating from the French domain (the nine clients associated with the French domain). The resulting performance is described in Figure 11. Compared to the previous results, the difference between the considered topology aggregations is exacerbated. More precisely, the portion of unfeasible demands accepted by the overprovisioned legacy aggregation increases from 1% to more than 10 % of all demands. On the other hand, the portion of erroneously rejected demands with under-provisioned aggregations increases from 10% to more than 30%, which is very significant. Also, for this load model, the second option of the under-provisioned aggregation exhibits better performance than the first one.

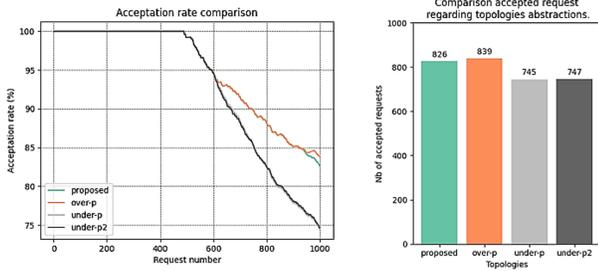


Figure 10: Service demand acceptance comparison for different Polish and French aggregated topologies.

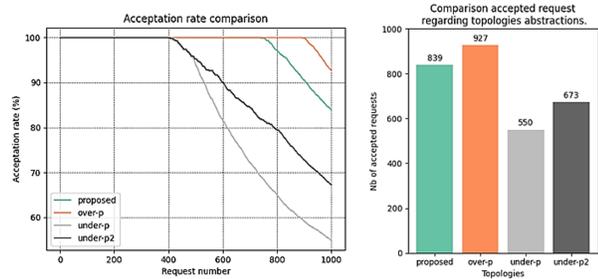


Figure 11: Demand acceptance for different aggregated topologies, with demands from the French domain

5.2 EVALUATION OF THE IMPACT OF DOMAIN-LEVEL SLICING

In this evaluation, we assume that some domains are implementing and exposing a domain-level slice providing point-to-multipoint (P2M) services within the boundaries of the domain and with some predefined quality of service. The rationale of exposing such slices is to allow, when possible, other domains to adopt these slices (i.e., assign slice resources from different domains) in order to build their end-to-end multi-domain P2M services, with the required QoS. We assume that if a domain exposes a P2M slice, this implicitly means that it considers that using the slice for an end-to-end P2M link is much more effective than relying on border node-to-border node

links. The objective of this evaluation is to compare the efficiency of the resource algorithm for P2M demands when P2M slices are exposed to the case where they are not.

5.2.1 Simulation set-up.

We consider randomly generated topologies with 25 domains and 15 sources and destinations randomly spread over the multi-domain network. Each domain is composed of 3 to 7 edge nodes. A random number of non-border nodes, ranging from 0 to 3, completes the network topology of each domain. At the domain and multi-domain level, we generate a random and consistent partial mesh topology. Service demands are composed of one single P2M end-to-end link with a randomly chosen number of destinations between 2 to 8. Their bandwidth and delay requirements are randomly generated as in the previous set of experiments. The number of service demands is varied without overloading the multi-domain network as the objective is to compare the efficiency of the two options when admitting the same demands. Also, when generating the random topologies, we vary the portion of domains that are exposing a P2M slice from 25% to 100%. Also, by default, we force the resource allocation algorithm to assign slice resources if the considered end-to-end link is duplicated when crossing the slice. This avoids the situation where slice resources are used to support point-to-point transmissions.

This is achieved by defining a new set of binary variables $s_i^k, \forall i \in N_s$ and $\forall k \in D$ which specifies whether slice i is being used by link k , more formally :

$$\forall k \in D, \forall i \in N_s, \forall (j, i) \in A' : s_i^k \geq x_{(j,i)}^k \quad (9)$$

The above-cited constraint is enforced as follows:

$$\forall k \in D, \forall i \in N_s : \sum_{(i,j) \in A'} x_{(i,j)}^k \geq 2 \cdot s_i^k \quad (10)$$

Finally, parameter α_1 of the objective function is set to 0 and the third term is given the highest weight. This means that, whenever possible, P2M slices are favored compared to edge-to-edge links.

5.2.2 Performance result.

The performance metric used to compare the efficiency of the resource allocation in the presence of exposed P2M slices is the percent decrease in the overall bandwidth that was allocated to an end-to-end link k when compared to the resources that were needed with no P2M slices. More precisely, when P2M slices are exposed, the overall amount of bandwidth assigned to an end-to-end link k is computed as follows:

$$\sum_{i \in N_s} \sum_{(j,i) \in A'} \phi_{j,i}^k + \sum_{\substack{(i,j) \in A' \\ i,j \notin N_s}} \phi_{i,j}^k \quad (11)$$

It is worth to note that we are only including ingress slices' links. With this solution, we implicitly assume that the P2M slices consume one single transmission to reach many destinations. This assumption holds for many physical networks, notably wireless networks, including satellite communication networks. Alternatively, the amount of resources could be computed by only considering slices' egress links, which is less favorable. Finally, when P2M slices are not exposed, the overall amount of assigned resources equals $\sum_{(i,j) \in A'} \phi_{i,j}^k$

Figure 12 and Figure 13 present the results when P2M slices are respectively available at 25% of the domains and 75%. Each point reflects the percent decrease of an end-to-end P2M link; links are classified according to the number of destinations (from 2 to 8 on the x-axis).

The first observation is that when the portion of P2M slices is low, many demands (which are randomly generated) do not take advantage of the presence of P2M slices because of the location of their endpoints. This is particularly the case of end-to-end links with 2-3 destinations and this is normal. Indeed, the greater the number of destinations is, the more the probability of crossing a nation well located which provides a P2M slice is important. When the number of destinations increases, many more links take advantage of P2M slices. However, an average gain around 2,3% is visible.

By increasing the portion of domains that host P2M slices, the number of links that benefit from P2M slices increases (Figure 13). Results are much more significant than in the previous case. For some links, the gain reaches 40% with the ingress calculation method. Even more interesting, the average benefit of traversing slices can reach 25% of gain with 8 destinations. As a general conclusion, exposing such efficient P2M slices may lead to high resource savings.

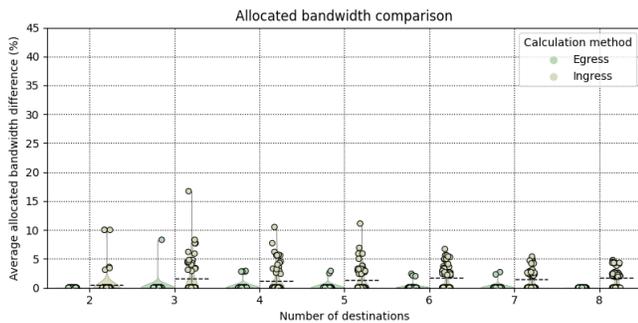


Figure 12: Allocation comparison (25% domains with P2M).

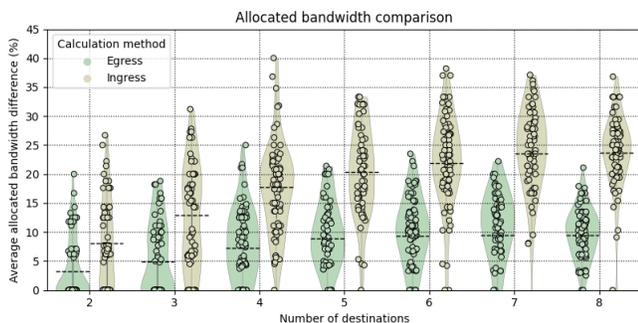


Figure 13: Allocation comparison (75% domains with P2M).

6 CONCLUSIONS

Deriving a legacy domain topology aggregation is not that trivial especially for complex and large domains. We have shown that even if the abstractions are derived with care, they may lead to a non-negligible degradation of demand admissibility (with an under-provisioned aggregation) or an increase of demand admission delays and network and computing overhead due to erroneously accepted demands at the multi-domain. This clearly justifies the topology aggregation that we propose, which allows the inclusion of abstracted non-border nodes in order to build more precise and effective abstractions. We have also shown that introducing appropriate domain-level slices into the topology aggregation of domains offers many advantages, notably a more efficient resource usage. In addition, it leads to simplified and more scalable resource allocation methods.

ACKNOWLEDGMENTS

This work was supported by the European Defence Agency (EDA) project No B 1520 IAP4 GP "Software Defined Tactical and Theatre Network (Softanet)".

REFERENCES

- [1] Ahmed Abujoda and Panagiotis Papadimitriou. 2016. DistNSE: Distributed network service embedding across multiple providers. In *2016 8th International Conference on Communication Systems and Networks (COMSNETS)*. 1–8. <https://doi.org/10.1109/COMSNETS.2016.7439948> ISSN: 2155-2509.
- [2] Zuse-Institute Berlin. 2006. SNDlib. <http://sndlib.zib.de/home.action>
- [3] Mikael Capelle, Slim Abdellatif, Marie-José Huguët, and Pascal Berthou. 2015. Online virtual links resource allocation in Software-Defined Networks. In *2015 IFIP Networking Conference (IFIP Networking)*. 1–9. <https://doi.org/10.1109/IFIPNetworking.2015.7145320>
- [4] David Dietrich, Ahmed Abujoda, Amr Rizk, and Panagiotis Papadimitriou. 2017. Multi-Provider Service Chain Embedding With Nestor. *IEEE Transactions on Network and Service Management* 14, 1 (March 2017), 91–105. <https://doi.org/10.1109/TNSM.2017.2654681> Conference Name: IEEE Transactions on Network and Service Management.
- [5] David Dietrich, Amr Rizk, and Panagiotis Papadimitriou. 2015. Multi-Provider Virtual Network Embedding With Limited Information Disclosure. *IEEE Transactions on Network and Service Management* 12, 2 (June 2015), 188–201. <https://doi.org/10.1109/TNSM.2015.2417652> Conference Name: IEEE Transactions on Network and Service Management.
- [6] Andreas Fischer, Juan Felipe Botero, Michael Till Beck, Hermann de Meer, and Xavier Hesselbach. 2013. Virtual Network Embedding: A Survey. *IEEE Communications Surveys & Tutorials* 15, 4 (2013), 1888–1906. <https://doi.org/10.1109/SURV.2013.013013.00155> Conference Name: IEEE Communications Surveys & Tutorials.
- [7] Geir Hallingstad and Sander Oudkerk. 2008. Protected core networking: an architectural approach to secure and flexible communications. *IEEE Communications Magazine* 46, 11 (Nov. 2008), 35–41. <https://doi.org/10.1109/MCOM.2008.4689242> Conference Name: IEEE Communications Magazine.
- [8] Ines Houidi, Wajdi Louati, Walid Ben Ameer, and Djamal Zeglache. 2011. Virtual network provisioning across multiple substrate networks. *Computer Networks* 55, 4 (March 2011), 1011–1023. <https://doi.org/10.1016/j.comnet.2010.12.011>
- [9] Meng Shen, Ke Xu, Kun Yang, and Hsiao-Hwa Chen. 2014. Towards efficient virtual network embedding across multiple network domains. In *2014 IEEE 22nd International Symposium of Quality of Service (IWQoS)*. 61–70. <https://doi.org/10.1109/IWQoS.2014.6914301> ISSN: 1548-615X.
- [10] Tarik Taleb, Ibrahim Afolabi, Konstantinos Samdanis, and Faqir Zarrar Yousaf. 2019. On Multi-Domain Network Slicing Orchestration Architecture and Federated Resource Control. *IEEE Network* 33, 5 (Sept. 2019), 242–252. <https://doi.org/10.1109/MNET.2018.1800267> Conference Name: IEEE Network.
- [11] Nassima Toumi, Djamal-Eddine Meddou, and Adlen Ksentini. 2019. A multi-objective SFC placement scheme over multiple domains. In *ICC 2019, IEEE International Conference on Communications-20-24 May 2019, Shanghai, China*, IEEE (Ed.). Shanghai.