



HAL
open science

Controllable virtual network service over a multiadministrative multi-domain network

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

► **To cite this version:**

Stanislas Pédebéarn, Slim Abdellatif, Pascal Berthou, Dariusz Nogalski, Dallal Belabed. Controllable virtual network service over a multiadministrative multi-domain network. 27th IEEE International Symposium On Real-Time Distributed Computing (ISORC 2024), May 2024, Carthage, Tunis, Tunisia. à paraître. hal-04502334v1

HAL Id: hal-04502334

<https://hal.science/hal-04502334v1>

Submitted on 13 Mar 2024 (v1), last revised 15 Apr 2024 (v2)

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.

Controllable virtual network service over a multi-administrative multi-domain network

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

Slim Abdellatif
LAAS-CNRS, Universite de Toulouse,
CNRS, INSA Toulouse, France
slim@laas.fr

Pascal Berthou
LAAS-CNRS, Universite 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 —

In this paper, we propose a controllable virtual network service that can be provided on a multi-administrative multi-domain network. A resource allocation algorithm is proposed to compute the resources from different domains that are needed to support the virtual network with the required QoS. An implementation of the service on an OpenFlow enabled multi-domain network is described. The service is then applied in the context of coalition military network to show its benefits and potential.

Keywords — virtual network embedding, SDN, multi-administrative multi-domain networks, slicing, QoS

I. INTRODUCTION

Among the benefits brought by SDN, network softwarization and network virtualization is the ability to enable a variety of services that better meet current and forthcoming user needs. Since the advent of SDN and thanks to the flow-based forwarding promoted by SDN and OpenFlow, many connectivity services with various QoS guarantees, possibly set up on-demand and adjusted dynamically, were investigated and implemented in operational networks. In addition, services that go beyond providing connectivity by including network capabilities/functions were also extensively investigated (i.e. Service chaining/VNF chaining [7]). Another kind of service that is not intended for a predefined packet flow is providing a virtual network to support all the traffic exchanged between connecting hosts, with some predefined performance. Such virtual network is the interconnection of virtual nodes (hosted on physical network nodes) via virtual links established on top of the network infrastructure. It is not new and was considered for some time now, especially in a single domain context (a datacenter with the concept of Network as A Service (NAAS) [3]). Also, it is a way of implement network slicing [4].

This work proposes an end-to-end multi-domain service, which provides to the user a fully controllable/programmable virtual network that spans multiple domains. This network is composed of virtual nodes optimally scattered throughout the multi-domain network and connected with virtual links that, in turn, can span multiple domains. A network controller/manager is also part of the service as well as its connecting links to the virtual nodes (see **Figure 1**). Indeed, the idea is to provide a virtual network, whose behavior can be programmed and adjusted at will. One typical example of the service is an OpenFlow based multi-domain virtual network with its OpenFlow controller. Depending on the network control applications running on the controller, different behaviors or management policies can be

programmed on the virtual network. They can even change over time. Also, thanks to OpenFlow, virtual nodes support a finer-grained packet forwarding than what a conventional network nodes can do (which typically apply destination based forwarding).

In this paper, we investigate the advantages of such a service, we propose an ILP formulation for the embedding of our proposed service in a multi-domain context. We also propose a prototype implementation of the service on a heterogeneous platform that combines network emulators and physical SDN/OpenFlow switches and describe its operation by enforcing a basic traffic management policy.




Figure 1. Illustration of the proposed service

This paper is organized as follows. Section II describes the proposed service and the considered multi-domain substrate network. Section III proceeds with the mathematical formulation of the embedding of the proposed service on the considered substrate network model. Section IV describes the implementation of the service and demonstrates its operation. Section V highlights its benefits and some of its applications and discusses the related work from literature. Section VI concludes the paper and draws up some perspectives of this work.

II. CONSIDERED SERVICE AND MUTI-DOMAIN SUBSTRATE NETWORK

A. Proposed service

The programmable virtual network service that we are proposing is classically composed of three different parts (presented with different colors in Figure 2 :

- the data network in charge of conveying user traffic that is exchanged between connecting user hosts. It is composed of a set of virtual nodes that can be hosted on

different physical nodes located at different domains, and a set of virtual links connecting in one hand virtual nodes to each other and user hosts to the virtual network. Each virtual link is characterized by a bandwidth and delay requirement. Virtual nodes are characterized with a maximum packet switching rate and a forwarding table size.

- The control network, which is composed of the virtual network controller (which can be hosted on any server/datacentre from any domain) and the virtual links that connect the controller to all virtual nodes. It is in charge of conveying control and monitoring information that enable virtual network programmability. If we consider again the case of an OpenFlow based virtual network, this control network supports all OpenFlow traffic.

Virtual control links are characterized with a bandwidth and delay requirements. In turn, the virtual network controller has processing (CPU) resource requirements expressed in terms. The needed resources (control link bandwidth & processing) are derived from the network control applications that state the virtual network behaviour.

- Optionally, a management network/links may also be part of the service in order to allow interactions with the controller from remote locations. In other words, the virtual network controller and some network control applications can be executed on separate servers (in user premisses). Again, some bandwidth and processing resource requirements are typically expressed.




FIGURE 2. SERVICE COMPONENTS

B. Considered multi-domain substrate network

When embedding an end-to-end multi-domain service with some Quality of Service (QoS) guarantees, the initiating domain must solicit multiple domains to support portions of the service. In order to identify the appropriate portions (exhibiting the appropriate performance guarantees), combine them to compose the requested service, the initiating domain relies on the aggregated topologies exposed by other domains (i.e., compact portrayal of the network topology with the available resources of each domain). In a multi-administrative multi-domain context, domains restrict the topology information that they disclose to other domains. For instance, when dealing with multi-domain connectivity services with QoS on a multi-administrative multi-domain network, the aggregated topology that domains expose to the others is limited to border nodes (nodes connected to other domains) fully or partially connected to each other with abstract links. Some QoS/performance information are

attached to these latter, typically maximum and available capacity, maximum transfer delay, maximum packet loss rate, etc. Some cost information can also be used to enforce some policy-based decisions.

To embed the proposed service, domains need to expose more than border nodes and abstract links. Indeed, they need to reveal some nodes that are able to host virtual nodes, the type of virtual nodes that they can host and the available capacity, which is used to assess whether sufficient resources are available to support new virtual nodes instances. Also, some compute nodes (with the available computing resources) need to be disclosed in order to be considered as a potential host for the virtual network controller and eventually some network control applications.

As network slicing is gaining momentum, some domains may support multiple domain-level slices (each providing predefined types of domain-level services) that can be exposed and made available to other domains in order to compose their end-to-end multi-domain services. In such cases, slices can also be disclosed as part of the aggregated topology of a domain with the characterization of the type of services provided by the slice and the available capacity. 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 between a set of nodes exposed by the domain.

Lastly, in this work, domains can also expose non-border nodes, which allow them to abstract/compact some of their infrastructure topology constructs. As shown in [5], the inclusion of these abstracted non-edge nodes (that we call transport nodes) allows more precise multi-domain abstractions, leading to significant improvement of service demand admissibility and decent decrease of service demand admission delays, and network and computing overhead.




FIGURE 3. ILLUSTRATION OF TOPOLOGY ABSTRACTION COMPONENTS

Figure 3 sums up the different constructs that can be used by a domain to build its aggregated topology that it wants to disclose to other domains. Obviously, domains have sovereignty on which construct they want to use and on the aggregated topology they want to expose, and also, on how they map the resources attached to the exposed topology to the physical resources at the domain level. They even have sovereignty on the multi-domain substrate network that they derive from the aggregated topologies exposed by domains,

which is used for embedding the virtual network service demands.

III. MATHEMATICAL FORMULATION

This section describes the mathematical formulation that we propose to solve the online resource allocation problem of the controllable multi-domain virtual networks on top of the substrate multi-domain network described in the previous section. The output of the allocation is the assignment of the physical nodes that will host each virtual node and each controller, as well as the set of data paths that will support each node-to-node logical link and each controller-to-node logical link. We assume that path-splitting is not allowed to support all above-cited logical links. Without any loss of generality, we assume that one single controller is in use, and virtual nodes cannot be hosted on the same physical node. Also, the objective of the resource allocation is to distribute fairly network traffic and use efficiently network resources exposed by all domains. The proposed resource allocation method can be easily adjusted and extended to consider different options and integrate additional constraints regarding the placement of nodes and the controller(s).

Below, the substrate network and virtual network demand models are described. Then, the variables and problem constraints are listed. Lastly, the considered objective function is defined.

A. Multi-domain substrate Network model

The multi-domain substrate model is the interconnection of the aggregated topologies exposed by all domains. It is modeled as a unidirectional graph $G = (N, A)$ where N is the set of vertices, which correspond to exposed nodes. $A \subseteq N \times N$ is the set of edges (i.e. the set of exposed abstract links) connecting exposed nodes to each other. $N = N_t \cup N_c \cup N_s$, where N_t, N_c, N_s are the set of transport nodes (border nodes and also abstracted non-border nodes), compute nodes and slice nodes. Indeed, domain level slices are also modeled as nodes; they are connected to the nodes that either send to or receive packets from the slice. The performance of the service (delay, etc.) provided by the slice is reported on the corresponding edge.

To each node $i \in N_t$, is associated available switching capacity TE_i , which is the current available number of entries of its forwarding table. The maximum size of node i flow table is denoted by TE_i^{max} . In addition, node i is characterized by a maximum forwarding rate FR_i^{max} , which in fact reflects the amount of node i switching capacity (resources) that a domain reserves to other domains (expressed in terms of the maximum overall (whatever the incoming interface) number of packets that the node i forwards, on behalf of the federation, per unit of time). The remaining switching capacity at time of arrival of a new demand is FR_i . Similarly, for each slice $i \in N_s$, we respectively denote as FR_i^{max} and FR_i the maximum and remaining capacity (in terms of packet rate) that the slice can handle. Last, for each link $(i, j) \in A$, we respectively denote by $\gamma_{i,j}^{max}$, γ_{ij} and $l_{i,j}$ as the maximum capacity of the link as initially made available by the corresponding domain to the other domains (i.e. before any resource assignment), the available capacity of (i, j) at time of service demand arrival and the max packet transfer delay along the link.

We denote as $N_v \subseteq N_t \cup N_c$ the set of substrate nodes that are able to host virtual nodes. A substrate node $i \in N_v$ can either be a transport node or also, a compute node since many network nodes (Openflow switch, etc.) can be implemented in software running on servers. Without loss of generality, we assume that network controllers can only be run on compute nodes, and that all compute nodes (i.e. N_c) can host any network controller. It is worth to note that, without loss of generality, in this formulation, we are considering that network controllers are supported by the resources available at the multi-domain network (and not user premises). Also, we assume that any controller can be hosted by any compute node.

Any compute node and any transport node $i \in N_v$ that is capable of hosting a virtual node is characterized by an available and maximum node capacity, respectively denoted as γ_i and γ_i^{max} expressed in compute units. Last, the set of types of supported virtual nodes/networks (i.e. Openflow, IP router, Ethernet bridges, etc.) is denoted as F_{vnet} . Each type $tp \in F_{vnet}$ is characterized by a unitary resource cost $\gamma^{tp,unit}$, which specifies the amount of resources (in compute units) needed per units of *kbps* of traffic. The set of substrate nodes that support virtual nodes of type tp is denoted as N_{tp} .

B. Demand Model

Each demand is composed of a set of multi-domain virtual networks D . Each virtual network $k \in D$ of type tp^k is a unidirectional graph $G^k = (V^k, E^k)$ where:

- The type tp^k of virtual network k , can be OpenFlow, IP, Ethernet bridge, etc. Without losing generality, we assume here that virtual network k is composed of homogeneous nodes, i.e. all of the same type.
- V^k is the set of virtual nodes composing virtual network k as well as the network controller (denoted as ctl^k). $\forall v \in V^k - \{ctl^k\}$, v is of type tp^k .
- ctl^k 's computing resources requirements $\gamma^{k,ctl}$ expressed in compute units. The user can even express in the demand, the set of compute nodes that are allowed to host the network controller. It denoted as $N_{ctl}^k \subseteq N_c$. $\gamma^{k,ctl}$ is specific to each virtual network and typically depend from the number of nodes under the control of the network controller and from the network Operating System (Network OS) services that the controller implements and to a less extent from the network control applications that control the virtual network.
- $E^k \subseteq V^k \times V^k$ the set of virtual links connecting virtual nodes to each other and the network controller to virtual nodes. Each virtual link $(v, v') \in E^k$ has a bandwidth requirement of $b^{k,(v,v')}$, and a maximum transfer delay of $\underline{L}^{k,(v,v')}$
- a maximum packet size of p^k .

C. Resource-related assignment variables

The output of the resource allocation algorithm is the set of substrate nodes that will host each virtual node of each virtual network request $k \in D$ (with the required resources) and, also, the set of data paths (routes with the bandwidth allocations at each supporting substrate link and the number

of forwarding table entries at each crossed node) that will support each virtual link composing virtual network k (i.e. those connecting virtual nodes but also those that will transport control traffic between the network controller and the nodes). Hence, we distinguish the following variables:

- $z_i^{k,v}$: is a Boolean variable, which specifies whether virtual node $v \in V^k$ is embedded in the substrate node $i \in N_v$, i.e. $z_i^{k,v} = 1$, if node i hosts virtual node v , 0 otherwise. Also, we denote as $z_i^{k,ctl}$ the binary variable that indicates whether node $i \in N_c$ hosts the controller;
- $\phi_{i,j}^{k,(v,v')}$: represents the bandwidth assigned to the packets of virtual link $(v, v') \in E^k$ that are flowing from node v to node v' at link $(i, j) \in A$;
- $x_{i,j}^{k,(v,v')}$: is a Boolean variable which reflects whether the flow of packets of virtual link $(v, v') \in E^k$ is supported by the substrate link $(i, j) \in A$ (i.e., $x_{i,j}^{k,(v,v')} = 0$ if $\phi_{i,j}^{k,(v,v')} = 0$; 1 otherwise. These variables will be constrained to prevent any flow-splitting, i.e. one single data-path is used to support virtual link (v, v') ;
- $TE_i^{k,(v,v')}$: specifies the number of entries that are installed in node $i \in N_t$ forwarding table to support virtual link (v, v') . Many options can be considered to express switching resource consumption and integrated in the mathematical formulation. Without losing generality, we consider that the number of table entries needed by virtual link (v, v') equals one entry if at least one node i port is either receiving or sending traffic from (v, v') . Thus, $TE_i^{k,(v,v')}$ is a Boolean variable derived as follow:

$$\forall i \in N_t, \forall (j, i) \in A : x_{j,i}^{k,(v,v')} \leq TE_i^{k,(v,v')} \quad (1.a)$$

$$\forall i \in N_t, \forall (i, j) \in A : x_{i,j}^{k,(v,v')} \leq TE_i^{k,(v,v')} \quad (1.b)$$

$$\forall i \in N_t, \quad (1.c)$$

$$TE_i^{k,(v,v')} \leq \sum_{(j,i) \in A} x_{j,i}^{k,(v,v')} + \sum_{(i,j) \in A} x_{i,j}^{k,(v,v')}$$

D. Problem Constraints

The constraints related to the placement of virtual nodes $v \in V^k$ are described in inequalities 1 to 6. Equations 2 force nodes from virtual network k to be only placed on substrate nodes that support virtual node type tp^k . Same logic holds for the network controller in equations 3. Equations 4 ensure the placement of each virtual node on one substrate node. Equation 5 does the same for network controllers.

$$\forall k \in D, \forall v \in V^k - \{ctl^k\}, \forall i \in N_v - N_{tp^k} : z_i^{k,v} = 0 \quad (2)$$

$$\forall k \in D, \forall i \in N_c - N_{ctl^k} : z_i^{k,ctl} = 0 \quad (3)$$

$$\forall k \in D, \forall v \in V^k - \{ctl^k\} : \sum_{i \in N_v} z_i^{k,v} = 1 \quad (4)$$

$$\forall k \in D : \sum_{i \in N_c} z_i^{k,ctl} = 1 \quad (5)$$

Inequalities 6 constrain the placement of the vertices of a virtual link. Different options can be specified and easily expressed. One option is to force their placement on different domains, or even to constrain their placement on specific domain. The constraint that we consider below is to force their placement on different substrate nodes.

$$\forall k \in D, \forall (v, v') \in E^k \text{ with } v \neq ctl^k, v' \neq ctl^k, \forall i \in N_v :$$

$$z_i^{k,v} + z_i^{k,v'} \leq 1 \quad (eq 6)$$

Note that, the previous constraints allow the embedding the network controller and a virtual node on the same compute node. Of course, this can be easily prevented by including ctl^k to inequalities 6.

The constraints related to bandwidth allocations are described in equations 7 to 10. Equations 7 represent the usual flow conservation constraints related to virtual link $(v, v') \in E^k$ that connects two virtual nodes.

$$\forall k \in D, \forall (v, v') \in E^k \text{ with } v \neq ctl^k, v' \neq ctl^k,$$

$$\forall i \in N_s \cup N_t \cup N_v :$$

$$\sum_{(i,j) \in A} \phi_{i,j}^{k,(v,v')} - \sum_{(j,i) \in A} \phi_{j,i}^{k,(v,v')} \quad (7.a)$$

$$= \begin{cases} b^{k,(v,v')} \cdot (z_i^{k,v} - z_i^{k,v'}), & \forall i \in N_v \\ 0, & \forall i \in N_s \cup N_t - N_v \end{cases}$$

The next equations state that no bandwidth allocation is possible on links connected to compute nodes that do not support virtual nodes.

$$\forall k \in D, \forall (v, v') \in E^k \text{ with } v \neq ctl^k \text{ and } v' \neq ctl^k,$$

$$\forall i \in N_c - N_v, \forall (i, j) \in A, \forall (j, i) \in A:$$

$$\phi_{i,j}^{k,(v,v')} = \phi_{j,i}^{k,(v,v')} = 0 \quad (7.b)$$

Equations 8 is the flow conservation constraints that concern the virtual links between the controller and associated virtual nodes.

$$\forall k \in D, \forall (ctl^k, v') \in E^k, \forall i \in N:$$

$$\sum_{(i,j) \in A} \phi_{i,j}^{k,(ctl^k,v')} - \sum_{(j,i) \in A} \phi_{j,i}^{k,(ctl^k,v')} =$$

$$\begin{cases} b^{k,(ctl^k,v')} \cdot (z_i^{k,ctl} - z_i^{k,v'}), & \forall i \in N_v \cap N_c \\ -b^{k,(ctl^k,v')} \cdot z_i^{k,v'}, & \forall i \in N_v - N_c \\ b^{k,(ctl^k,v')} \cdot z_i^{k,ctl}, & \forall i \in N_c - N_v \\ 0, & \forall i \in N_s \cup N_t - N_v - N_c \end{cases} \quad (8a)$$

$$\forall k \in D, \forall (v, ctl) \in E^k, \forall i \in N:$$

$$\sum_{(i,j) \in A} \phi_{i,j}^{k,(v,ctl)} - \sum_{(j,i) \in A} \phi_{j,i}^{k,(v,ctl)} =$$

$$\begin{cases} b^{k,(v,ctl)} \cdot (z_i^{k,v} - z_i^{k,ctl}), & \forall i \in N_v \cap N_c \\ -b^{k,(v,ctl)} \cdot z_i^{k,ctl}, & \forall i \in N_c - N_v \\ b^{k,(v,ctl)} \cdot z_i^{k,v}, & \forall i \in N_v - N_c \\ 0, & \forall i \in N_s \cup N_t - N_v - N_c \end{cases} \quad (8b)$$

Equations 9 and inequalities 10 ensure that path splitting is not used to support virtual links. Equation 9 connects $\phi_{i,j}^{k,(v,v')}$ and $x_{i,j}^{k,(v,v')}$. Clearly, since path splitting is disabled,

the two variables are equivalent and one of them can be removed. This said, both variables are kept for the sake of clarity and generality and, more specifically, to easily derive the variant of this method that allows path splitting.

$\forall k \in D, \forall (v, v') \in E^k, \forall (i, j) \in A$:

$$\phi_{i,j}^{k,(v,v')} = b^{k,(v,v')} \cdot x_{i,j}^{k,(v,v')} \quad (9)$$

$\forall k \in D, \forall (v, v') \in E^k$:

$$\sum_{(i,j) \in A} x_{i,j}^{k,(v,v')} \leq 1 \quad (10a)$$

$$\sum_{(j,i) \in A} x_{j,i}^{k,(v,v')} \leq 1 \quad (10.b)$$

Inequalities 11 explicitly connect the placement of a virtual network element with the placement of the virtual links to which it belongs.

$\forall k \in D, \forall (v, v') \in E^k, \forall i \in N_v, \forall (i, j) \in A$:

$$z_i^{k,v} \geq x_{i,j}^{k,(v,v')} \quad (11.a)$$

$\forall k \in D, \forall (v', v) \in E^k, \forall i \in N_v, \forall (j, i) \in A$:

$$z_i^{k,v} \geq x_{j,i}^{k,(v',v)} \quad (11.b)$$

$\forall k \in D, \forall v \in V^k, \forall i \in N_v$:

$$z_i^{k,v} \leq \sum_{(v,v') \in E^k} \sum_{(i,j) \in A} x_{i,j}^{k,(v,v')} + \sum_{(v',v) \in E^k} \sum_{(j,i) \in A} x_{j,i}^{k,(v',v)} \quad (11.c)$$

The following inequalities ensure that the resource assigned to demand D do not exceed the remaining resources at the substrate network, namely substrate links, nodes, slices, and compute nodes.

Inequalities 12 ensure that the bandwidth assigned to all virtual links from all virtual networks of D does not exceed the remaining bandwidth on each substrate network's link.

$$\forall (i, j) \in A: \sum_{k \in D} \sum_{(v,v') \in E^k} \phi_{i,j}^{k,(v,v')} \leq \gamma_{ij} \quad (12)$$

Inequalities 13 ensure that bandwidth allocations assigned to all virtual links respect substrate nodes' and exposed slice's remaining forwarding capacities

$$\forall i \in N_t \cup N_s: \sum_{k \in D} \sum_{(v,v') \in E^k} \sum_{(j,i) \in A} \phi_{j,i}^{k,(v,v')} \leq FR_i \quad (13)$$

The constraints 14 ensure that the flow table entries needed by the virtual links composing the demand D remain below the remaining forwarding table size. As a variant to this formulation, if virtual nodes have as requirement a forwarding table size of, let's say, $TE^{k,v}$, then such requirement can be easily integrated in inequalities 14 by adding for nodes $i \in N_v$ the required table size $\sum_{k \in D} \sum_{i \in N_v} z_i^{k,v} \cdot TE^{k,v}$.

$$\forall i \in N_t: \sum_{k \in D} \sum_{(v,v') \in E^k} TE_i^{k,(v,v')} \leq TE_i \quad (14)$$

Inequalities 15 ensure that the remaining computing capacity in compute nodes and transport nodes suffices to support assigned virtual nodes and eventually network controllers. The first sum computes the resources consumed by the virtual node v . It adds the bandwidth from all virtual links that end at v weighted by the unitary resource cost (i.e. $\gamma^{tp^k,unit}$). Indeed, as stated above the resources consumed by a virtual node is proportional to the rate of incoming traffic. Referring to inequalities (2), if a node $\forall i \in N_c - N_v, z_i^{k,v} = 0$, this first term is zeroed.

$\forall i \in N_v \cup N_c$:

$$\sum_{k \in D} \sum_{v \in V^k - \{ctl^k\}} \sum_{\substack{(v',v) \in E^k \\ v' \in V^k - \{ctl^k\}}} z_i^{k,v} \cdot b^{k,(v',v)} \cdot \gamma^{tp^k,unit} + \sum_{k \in D} z_i^{k,ctl} \cdot \zeta^{k,ctl} \leq \gamma_i \quad (15)$$

The second term concern compute nodes that can host network controllers (refer to equalities (3)) and adds the amount of resources that controllers consume.

Inequalities 16 ensure that the maximum latency associated to each virtual link is met. As path-splitting is disabled, the transfer delay that a packet 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 $\frac{p^k}{b^{k,(v,v')}}}$ for packets belonging to virtual link (v, v') and switching and propagation delays bounded by latency $l_{i,j}$ for any link $(i, j) \in A_t \cup A_s$, where A_t (resp. A_s) are the set of edges in A whose one of their end is either a transport node or a slice node. The latency constraints are- as follows.

$\forall k \in D, \forall (v, v') \in E^k$:

$$\sum_{(i,j) \in A_t \cup A_s} x_{i,j}^{k,(v,v')} \cdot \left(\frac{p^k}{b^{k,(v,v')}} + l_{i,j} \right) \leq L^{k,(v,v')} \quad (16)$$

E. Objective function

Different objective functions can be adopted with this formulation. Basically, the set of variables used in this formulation allow expressing different ways of efficiently using the substrate network resources and also effectively spreading the assigned resources in order to improve the admissibility/acceptance of forthcoming requests. One option is to minimize the highest link/node utilization rate of the substrate network. Another option is to minimizing a utilization threshold that apply to all substrate network elements. Another criterion is to minimize the number of active elements (compute nodes, transport nodes and even active links). Of course, all these objectives can be expressed with this formulation. Below, the objective function of equation (17) is only focused on minimizing the node and link resources that are needed to support the request D . Four terms are composing the objective function. The first aim at minimizing the average utilization of network substrate's links while the second aim to minimize the average utilization

of substrate nodes' forwarding table size. The objective of the last two terms is to minimize the average computing resources utilization that are used to support virtual nodes and network controllers on substrate network's nodes. Other similar sub-objectives can also be added, especially those related to the packet forwarding capacity of transport and slice substrate nodes.

Minimize

$$\begin{aligned}
& \alpha_1 \cdot \frac{1}{|A''|} \cdot \sum_{(i,j) \in A''} \left(\frac{1}{\gamma_{i,j}^{max}} \right. \\
& \quad \left. * \left(\gamma_{i,j}^{max} - \gamma_{i,j} + \sum_{k \in D} \sum_{(v,v') \in E^k} \phi_{i,j}^{k,(v,v')} \right) \right) \\
& + \alpha_2 \cdot \frac{1}{|N''_t|} \cdot \sum_{i \in N_t} \left(\frac{1}{TE_i^{max}} \right. \\
& \quad \left. * \left(TE_i^{max} - TE_i \right. \right. \\
& \quad \left. \left. + \sum_{k \in D} \sum_{(v,v') \in E^k} TE_i^{k,(v,v')} \right) \right) \\
(17) \quad & + \alpha_3 \cdot \frac{1}{|N''_v|} \cdot \sum_{i \in N_v} \left(\frac{1}{\gamma_i^{max}} \left(\gamma_i^{max} - \gamma_i \right. \right. \\
& \quad \left. \left. + \sum_{k \in D} \sum_{v \in V^{k-\{ctl^k\}}} z_i^{k,v} \cdot b^{k,(v',v)} \cdot \gamma_{tp^k,unit} \right) \right) \\
& + \alpha_4 \cdot \frac{1}{|N''_c|} \cdot \sum_{i \in N_c} \left(\frac{1}{\gamma_i^{max}} \left(\gamma_i^{max} - \gamma_i + \sum_{k \in D} z_i^{k,ctl} \cdot \zeta^{k,ctl} \right) \right)
\end{aligned}$$

IV. SERVICE IMPLEMENTATION & RESULTS

A. Considered use case & objectives

We consider federated military networks [1] to showcase our proposed service. In these networks ally nations connect and share some of their network infrastructure to build a multi-domain Federated Mission Network (FMN) that span a large region with points of presence at the theater of operation but also at very remote headquarters or command posts or datacenters. The goal is to leverage on this multi-domain FMN to provide the end-to-end multi-domain communication services needed during the coalition military missions [1]. The end-to-end service provided on such network is standard IP connectivity between end-points. Also, a crucial aspect is to efficiently use the resources provided by nations and maximize the support of the end-to-end services needed for the mission.

In this context, the service that we are proposing allows a nation to deploy and operate wide area virtual networks whose nodes can be located at very remote locations and

provided by different nations. Since the requesting nation is granted full control to its virtual networks, it can enforce its own policies (QoS, security, etc.) regardless of the policies used by the other nations that contribute to the provision of the virtual network. Moreover, these policies can be changed over time depending on mission's needs. Also, each virtual network can be programmed to provide customized and novel services (not only IP). This is particularly the case when provisioning an OpenFlow based Virtual network, which allows a flow-based forwarding at the level of the virtual network (as opposed to standard IP destination-based forwarding).

The objectives of the implementation are twofold. First, to show the feasibility of provisioning an OpenFlow-based multi-domain virtual network on an OpenFlow-enabled multi-domain substrate network. Second, demonstrate the operation of the proposed service by enforcing an ad-hoc QoS policy and highlight some of its benefits.

B. Implemented prototype

The network platform is composed of three domains belonging to three different nations. Figure 4 describes the considered multi-domain substrate network combining the aggregated topologies exposed by the three domains. For simplicity, we assume that the exposed topologies resume to domain level topologies. This avoids dealing with the topology aggregation policy of each domain. Hence, it avoids the need of having a domain level resource allocation since the embedding that we propose at the multi-domain (i.e., aggregated) level embeds the virtual network by considering all domain-level resources (and not an abstraction of these). Each domain is an OpenFlow-based network with one domain level OpenFlow controller and a couple of OpenFlow switches. OpenFlow traffic between the controller and the switches are exchanged in "in-band" mode by assigning a dedicated VLAN and installing the appropriate OpenFlow rules at each switch. This is part of the domain network initialisation. Most switches are emulated OVS (OpenVswitch) switches with the Containernet emulation tool below [6]. These are depicted in green in the figure. Some others are PICA8 p3295 physical switches that, in turn, implement a hardware optimized version of OVS. We assume that only the physical switches have the capability to host OpenFlow virtual switches. A datacentre is also exposed by domain 3. From a network management perspective, each domain has a "Service Orchestrator" (SO) in charge of orchestrating the provisioning of the multi-domain services initiated by the domain. When solicited by other domains, it also contributes to the provisioning of some portions of their accepted services [8]. For simplicity, the management traffic exchanged between the SOs is supported "out-of-band" via a dedicated management network. Each domain has a dedicated computer to run its Containernet network emulation tool, the Ryu controller and the Service Orchestrator.

The resource allocation algorithm was implemented on the Cplex solver. It is executed by the "Service Orchestrator" of the initiating domain upon the arrival of the service demand. Under the above-cited assumption, when the demand is accepted, the algorithm outputs all the required resources

(related links, switches, compute nodes from all crossed domains) for each virtual network. The SO launches the service deployment stage by requesting the deployment of virtual nodes, virtual links, and controllers. On its domain, it may instruct its physical switches to instantiate one or multiple virtual nodes, its compute nodes to instantiate a controller, and its domain's OpenFlow controller to generate the appropriate OpenFlow rules to set-up some portions of virtual links. Also, it may instruct other SOs to provision the needed resources (virtual nodes, portions virtual links, etc.) on the domain they belong to.

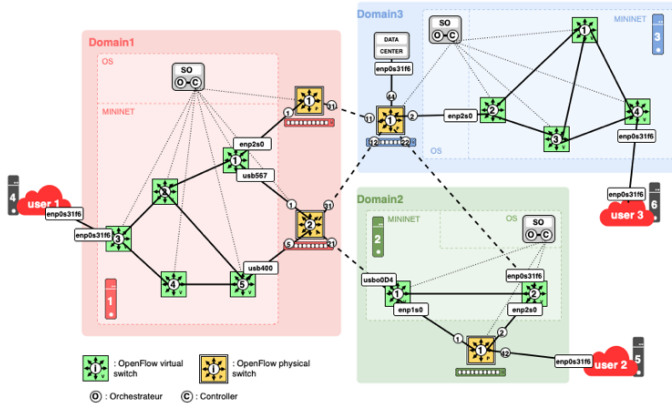


FIGURE 4 NETWORK SET-UP

For simplicity, each virtual link is assigned a unique VLAN identifier within a domain. To provision a virtual link, OpenFlow rules that match this VLAN identifier are installed on the switches along its path. If the virtual link crosses another domain, the VLAN identifier is rewritten accordingly. Clearly, this solution is not scalable with the number of virtual links and can only be considered for prototyping. A more effective solution is to rely on MPLS (Multi-protocol label Switching) labels or even on appropriately formatted MAC addresses.

Regarding virtual OpenFlow node instantiation, they take place at the physical OpenFlow (OF) switches, which in fact implement OVS. A first OF bridge is instantiated, and all the physical links to which the physical switch is connected to are attached to it. We denote this switch as OF-sw. The instantiation of a virtual OF node is realized by instantiating a new OF bridge. We denote this virtual node as OF-vnode. As shown in Figure 5, OF-vnode is connected to OF-sw via multiple “patch” ports: One “patch” port per virtual data link and one “patch” port for the control link that conveys OpenFlow traffic to its controller. As each virtual link is identified with a specific VLAN, “patch” ports are attached to their VLAN and OpenFlow rules based on the appropriate VLAN-ID are installed on OF-sw to route the data and control traffic to the right virtual node.




FIGURE 5 VIRTUAL NODE IMPLEMENTATION

C. Results Overview

1) Service request provisioning

We consider below a service request composed of one OF-based virtual network presented in Figure 6. This latter specifies a topology which represents the expected virtual network with some QoS requirements (bandwidth, latency, etc.). In our case, we are deploying a virtual network composed of 3 OpenFlow virtual switches and one network controller. Three user networks are reachable from this virtual network.




Figure 6 service request

Figure 7 describes the resources computed by the resource allocation algorithm and assigned to the requested service across the multi-domain network. The data network component of the virtual network is as follows. The virtual nodes were placed on three different physical switches (right-hand side) as enforced by the embedding algorithm. The embedding virtual data links is depicted in blue. The controller of the virtual network is placed on the unique compute node belonging to domain 3. The three virtual control links, which also compose the control component of the virtual network, are depicted in green. Last, the embedding of the management link that connect the controller to user1's network is depicted in red.

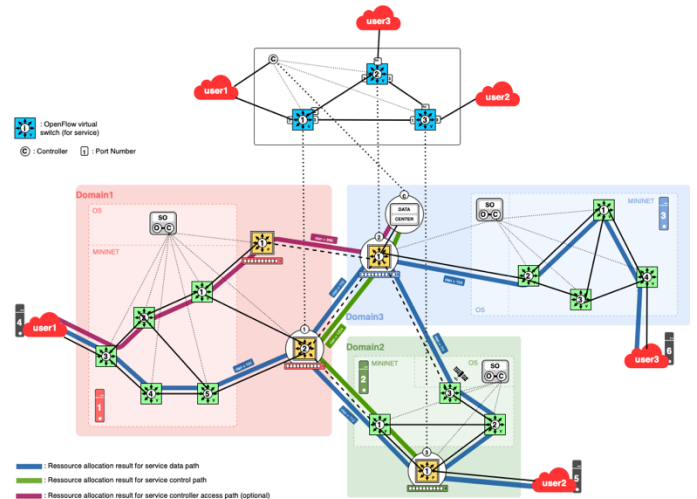


FIGURE 7 PROVISIONED RESOURCES

2) Virtual network operation

To show the operation of this service, the virtual network of **Figure 8** is requested and deployed with a dedicated control application. The control application demonstrates the ability to enforce any ad-hoc QoS policy on the virtual network. This policy is a basic representation of the policies used in military networks. Two levels of priorities are considered: one for video streams and one for low-priority data streams. The QoS policy implemented in the control application ensures that at any moment a high priority flow (video stream) has precedence on data streams on any network resource. When a new video stream is launched, if resources are available (not used by already admitted video flows), the control application programs the virtual switches to route the video stream. In order to protect the resources assigned to the new video stream, when possible, some existing data streams may be rerouted along longer and less efficient paths. If such paths are not available, the data stream is interrupted, i.e., its traffic is filtered/dropped at network entrance. **Figure 9** highlights the control application workflow.

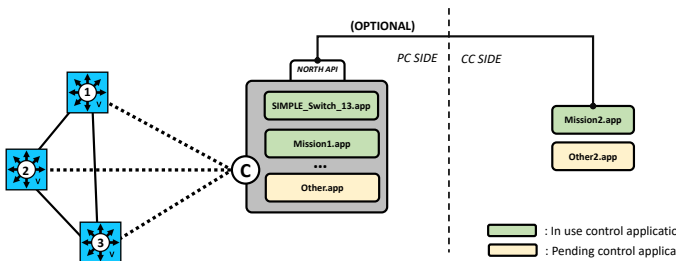


Figure 8 - Service request with control application

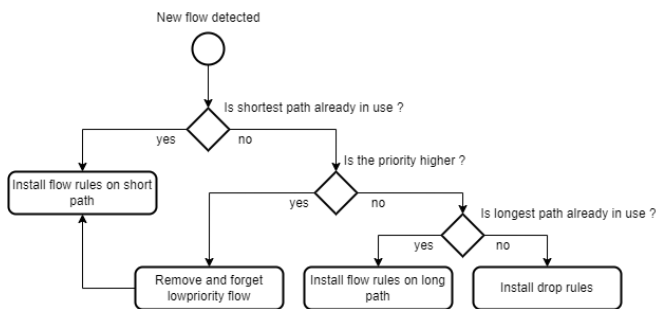


Figure 9: Demonstration dedicated control application working flow

The following scenario, played to illustrates the enforcement of the QoS policy on the deployed virtual network. First a data stream is detected from user2 to user1 and reported to the network controller, which assigns to the flow the shortest path (**Figure 10**, pin 1). A video stream is set from user2 to user1, the controller detects a resource constraint violation on the shortest path. Then it assigns the shortest path to the newly arrived video flow and reroute the data stream on the longer path via the upper part of the network via virtual switch 2 (**Figure 10**, pin 2). A video stream is the set from user3 to user1, no more free paths are available, the controller decides to assign the shortest path to the video stream and to drop the data stream (**Figure 10**, pin 3). When, the second video stream stops, the network controller detects it and re-assigns the data-path to the data stream. This latter restarts on the long path. When, the first video stops, the controller

detects it and the data stream is redirected to the shortest path. All of the logic is coded in the network control application run at the controller.




Figure 10: Main steps of the considered scenario

V. DISCUSSION

The service proposed in this paper offers to the user the opportunity to have a dedicated network that span multiple administrative domains that it can fully manage and fully control in order to provide customized multi-domain services. This network is composed of virtual nodes optimally scattered throughout the multi-domain network. When using Openflow, virtual nodes support a finer-grained packet forwarding than what an IP router can do. Sophisticated, homogeneous and dynamic traffic management policies can be enforced on the virtual network even if virtual nodes are hosted in different domains. Also, different control applications can be executed by the controller over time, allowing the user to adjust the forwarding policy according to its current needs. Naturally, many such virtual networks can be provisioned at the same time through the multi-domain network. They can be requested on demand and, possibly, deployed in the time scale of minutes rather than days.

We have shown that our proposal can be useful in a FMN context. We also think similar interest exists in the civilian context since service providers usually have privileged partners; their respective domains may enable such service for their own use or for their clients.

This service assumes that, based on network/device virtualization, domains will “logically” share their network devices with other domains and provide them with some decent programming capabilities on these virtualized devices. This is much more demanding and committing than providing “best-effort” IP transit, as in the current multi-domain networks, and puts a lot of security and performance commitments/requirements on the hypervisor in charge of device virtualization.

Virtual network embedding has been extensively researched for more than a decade [9],[10],[11]. Different variants of the algorithms were proposed by considering various resources, resource cost and optimization criteria. To the best of our knowledge, the peculiarities of our proposed variant are as follows. First, the consideration of the three components of the virtual network: data, control and management components. Also, our algorithm considers the switching resources: forwarding table size as well as the forwarding rate. Last, it takes into account domain-level network slices.

VI. CONCLUSION

A controllable virtual network deployed on a multi-administrative multi-domain network offers interesting opportunities. The ability to enforce homogeneous and

dynamic traffic management policies on the virtual network even if the virtual nodes are hosted by different domains. The ability to leverage the virtual network to provide customized and novel end-to-end service on top of a multi-domain network. In this paper, we have proposed a resource allocation algorithm for such service. We have also described a prototype implementation on top of an SDN enabled multi-domain network and elaborated on its benefits.

Some perspectives to this work are as follows. To consider the case of dynamic virtual network with a changing topology or QoS requirements. Another point is how to derive the virtual network topology and performance requirements from high-level user intents.

ACKNOWLEDGMENT

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

REFERENCES

- [1] Hallingstad, Geir, and S. Oudkerk. "Protected Core Networking: An Architectural Approach to Secure and Flexible Communications". *IEEE Communications Magazine* 46, no. 11 (Nov. 2008): 35–41.
- [2] A. Fischer, J. F. Botero, M. T. Beck, H. de Meer and X. Hesselbach, "Virtual Network Embedding: A Survey," in *IEEE Communications Surveys & Tutorials*, vol. 15, no. 4, pp. 1888-1906, 2013.
- [3] A. Boubendir, E. Bertin and N. Simoni, "Flexibility and dynamicity for open network-as-a-service: From VNF and architecture modeling to deployment," *NOMS 2018 - 2018 IEEE/IFIP Network Operations and Management Symposium*, Taipei, Taiwan, 2018, pp. 1-6, doi: 10.1109/NOMS.2018.8406135.
- [4] L. U. Khan, I. Yaqoob, N. H. Tran, Z. Han and C. S. Hong, "Network Slicing: Recent Advances, Taxonomy, Requirements, and Open Research Challenges," in *IEEE Access*, vol. 8, pp. 36009-36028, 2020, doi: 10.1109/ACCESS.2020.2975072
- [5] Stanislas Pedebearn, Slim Abdellatif, Pascal Berthou, Dariusz Nogalski, Dallal Belabed, Virtual Link Embedding in Collaborative Sliced Multi-administrative Multi-Domain Networks, 39th ACM/SIGAPP Symposium On Applied Computing, Avila, Spain, April 8 - April 12, 2024
- [6] Containernet, <https://containernet.github.io>, consulted 01/2024
- [7] Michael Till Beck, Juan Felipe Botero: Scalable and coordinated allocation of service function chains. *Comput. Commun.* 102: 78-88 (2017)
- [8] Nathan F. Saraiva de Sousa, Danny A. Lachos Perez, Raphael V. Rosa, Mateus A.S. Santos, Christian Esteve Rothenberg, Network Service Orchestration: A survey, *Computer Communications*, Volumes 142–143, 2019, Pages 69-94
- [9] H. Cao, S. Wu, Y. Hu, Y. Liu and L. Yang, "A survey of embedding algorithm for virtual network embedding," in *China Communications*, vol. 16, no. 12, pp. 1-33, Dec. 2019, doi: 10.23919/JCC.2019.12.001.
- [10] Stefan Tomaszek, Lars Fritsche, Andy Schürr, "Dynamic Virtual Network Embedding: Using Incremental Model Transformation and Integer Linear Programming Techniques", *Journal of Object Technology*, Volume 19, no. 2 (July 2020), pp. 9:1-21, doi:10.5381/jot.2020.19.2.a9
- [11] A. Fischer, J. F. Botero, M. T. Beck, H. de Meer and X. Hesselbach, "Virtual Network Embedding: A Survey," in *IEEE Communications Surveys & Tutorials*, vol. 15, no. 4, pp. 1888-1906, Fourth Quarter 2013, doi: 10.1109/SURV.2013.013013.00155
- [12] M. Ait aba, M. Elkael, B. Jouaber, H. Castel-Taleb, A. Araldo and D. Olivier, "A two-stage algorithm for the Virtual Network Embedding problem," *2021 IEEE 46th Conference on Local Computer Networks (LCN)*, Edmonton, AB, Canada, 2021, pp. 395-398, doi: 10.1109/LCN52139.2021.9524968.