High-Performance Routing for Hose-Based VPNs in Multi-Domain Backbone Networks

Xiuzhong Chen[a,c,*], Marc De Leenheer[b,c], Rui Wang[c], Chaitanya S.K. Vadrevu[c], Lei Shi[c], Jie Zhang[a], Biswanath Mukherjee[c]

[a]*Key Lab. of Information Photonics and Optical Comm. (Beijing University of Posts & Telecom.), MOE, China*
[b]*Ghent University - IBBT, Belgium*
[c]*University of California - Davis, USA*

Abstract

By utilizing Layer-1 Virtual Private Networks (L1VPN), a single physical network, e.g., an optical backbone network, can support multiple virtual networks, which form the basic infrastructure for cloud computing and other enterprise networks. The L1VPN hose model is an elegant and flexible way to specify the customers' bandwidth requirements, by defining the total incoming and outgoing demand for each endpoint. Furthermore, multi-domain physical infrastructures are common in L1VPNs, since these are usually deployed on a global scale. Thus, high-performance Routing for Multi-domain VPN Provisioning (RMVP) for the hose model is an important problem to efficiently support a global virtual infrastructure. In this paper, we formulate the RMVP problem as a Mixed Integer Linear Program (MILP). Also, we propose a Top-Down Routing (TDR) strategy to compute the optimal routing for the hose model L1VPN in a multi-domain backbone network. Results indicate that TDR approaches the minimum routing cost when compared to the ideal case of single-domain routing.

Keywords: high-performance routing, layer-1 virtual private network, hose model, multi-domain, virtual infrastructure

---

[*]Corresponding author
Email address: `xzchen@ucdavis.edu` (Xiuzhong Chen)

1. Introduction

Layer-1 Virtual Private Networks (L1VPN) are logical networks that are established on a shared physical infrastructure, e.g., global optical backbone networks, to obtain a private communication environment without investment in physical network resources [1]. For an Infrastructure Provider (InP), L1VPN can improve the utilization of the physical network and increase profit. In particular, both public and private cloud computing platforms can be based on virtual infrastructure provisioned over a global network [2].

Two popular models exist to describe the traffic demands in a L1VPN, based on what information is available on the bandwidth demands of the Customer Edge (CE) devices[1]. One is the *pipe model*, in which the bandwidth demand for each pair of CEs is given, i.e., the traffic matrix of the pipe model is fixed, as shown in Fig. 1(a). An alternative is the *hose model*, in which the total incoming and outgoing bandwidth demand for each individual CE is given, as shown in Fig. 1(b). Consequently, the key feature of the hose model is that the traffic demand for any CE pair can be variable during its lifetime. This allows more flexibility to the CEs [3], although it has been difficult to create a high-performance routing algorithm for the provisioning of L1VPNs under the hose model. One solution that has recently been proposed [4] will be discussed and extended in the paper.
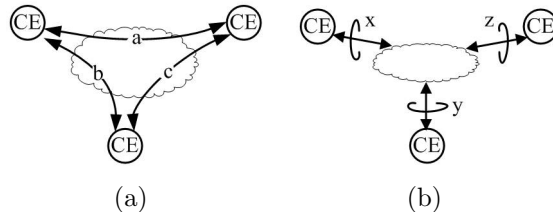

Figure 1: Traffic demands in L1VPNs: (a) Pipe Model and (b) Hose Model.

As L1VPNs are generally deployed on a worldwide scale, they cannot be provisioned using only a single InP. Global L1VPNs are composed of multi-domain networks, where the details within each domain (intra-domain) usually remain hidden from the other domains [5]. Thus, solutions to high-performance Routing for Multi-domain VPN Provisiong (RMVP) for global

---

[1]Customer Edge (CE) devices are the points in the network where traffic originates or arrives, and can thus be considered the end users of the network.

L1VPN services are desirable but hard to obtain. In this paper, we formulate the RMVP problem under the hose model by a Mixed Integer Linear Program (MILP) and propose an efficient Top-Down Routing (TDR) strategy for RMVP. One interesting global L1VPN scenario is designed to investigate our formulation and the TDR approach.

The rest of this paper is organized as follows. Related work is presented in Section 2. The RMVP model and TDR strategy are described in Sections 3 and 4. Numerical results are given in Section 5, and Section 6 concludes this paper.

2. Related Work

Provisioning under the pipe model has been well studied [6]. The hose model, introduced recently [3], is starting to receive attention [7, 8, 4]. Specifically, tree routing was studied where the network links have infinite capacity [7]. The authors presented a polynomial algorithm to compute the optimal tree routing in case of symmetric incoming and outgoing bandwidth demands, and proved it NP-hard to compute the optimal tree routing for general traffic demands (which may be asymmetric). In [8], the author improved the algorithm for computation of tree routings for general traffic demands and infinite link capacities. After that, the first polynomial-size LP formulation for maximum throughput routing of hose traffic demands along direct source-destination paths was developed by using the duality of Linear Program in [4].

Some other works studied the bandwidth efficiency of the hose model. A comparison of the hose and pipe models was done in [9], by defining the overprovisioning factor, which represents the additional capacity required for a VPN reservation in the hose model compared to the pipe model. In [10], the authors investigated the efficiency of the hose model by comparing tree routing, single-path routing, and multi-path routing, and demonstrated that multi-path routing offers significant advantages compared to the other two approaches.

A related problem is Virtual Network (VN) embedding, where the mapping of both virtual nodes and links is performed. In [11], the authors proposed VN embedding algorithms with coordination of nodes and links mapping. Then, the authors presented a policy-based VN embedding across multiple domains that embeds end-to-end VNs in a decentralized manner

in [12]. However, these works focus solely on the problem under the pipe-model VPN, and are difficult to directly extend to the hose model.

To the best of our knowledge, Routing for Multi-domain VPN Provisioning (RMVP) under the hose model has not been solved in the literature before. In this work, we formulate RMVP with hose model as a MILP, and propose a Top-Down Routing (TDR) strategy to solve the RMVP problem.

3. Problem Formulation

We outline the RMVP problem under the hose model in backbone networks which have mesh topology. We wish to minimize the cost of the bandwidth reserved for a L1VPN.

Let V_p denote the set (size N_p) of InP network nodes and V_x denote the set (size N_x) of nodes in domain x (thus, $V_x \subseteq V_p$) which belong to domain set D. Furthermore, let V_c denote the set (size N_c) of L1VPN Customer Edge (CE) devices, which connect with a single Provider's Edge (PE) device which are in turn elements of V_p. As such, the number of nodes in the complete topology is $N = N_c + N_p$. E is the set of links and contains inter-domain links, intra-domain links, and the links connecting CEs to their corresponding PEs. Inter-domain links are formed by a pair of nodes with each node in a different domain, and are denoted by L_{ij}. On the other hand, both nodes of an intra-domain are situated in the same domain, and are denoted as l_{ij}. Each link, both inter-domain and intra-domain, has an associated cost-per-unit-capacity value of Co_{ij} and a capacity C_{ij}. The hose-based traffic demand from CE m is defined by B_m^- (respectively B_m^+), indicating the incoming (respectively outgoing) bandwidth demand. The actual traffic matrix, composed of the bandwidth demands between any pair of CEs m and n, is given by t_{mn}. Note that, in the hose model, the traffic matrix is not known in advance, so multiple matrices may be feasible, making it very flexible.

To solve the RMVP problem, first consider an ideal case where all intra-domain topologies are known, which reduces the problem to minimization of the total bandwidth cost in the single domain. Any solution for the single-domain L1VPN forms a lower bound on the RMVP problem, since the intra-domain topologies are not known in a multi-domain routing scenario.

For a single domain, we define the node set as $V = V_p \cup V_c$, which contains both V_p and V_c. The objective is to minimize the total cost of bandwidth reserved for L1VPN demands. Since the feasible traffic matrix of hose model may be variable, the reserved bandwidth for each L1VPN should be sufficient

4

for the *worst case* of all traffic matrices. For illustration, consider Fig. 2 where three CEs are attached to a 7-node network. The traffic demands of the CEs are $B_m^- = B_m^+ = 8, \forall m \in \{A, B, C\}$. We assume that routing for this VPN is done over disjoint paths for each CE pair, e.g., t_{AB} passes through nodes 1, 2, 3, while t_{AC} passes through nodes 1, 6, 5. All other routing paths are shown as red arrows in Fig. 2. If the traffic matrix between any two nodes m and n equals $t_{mn} = 4, \forall m \neq n, \forall m, n \in \{A, B, C\}$, the hose demands can be routed by reserving a bandwidth of 4 on each link. However, in case the traffic matrix is changed to $t_{AB} = 8, t_{AC} = 0$ (which also satisfies the given hose demand), the reserved bandwidths for links l_{12} and l_{23} are 8.
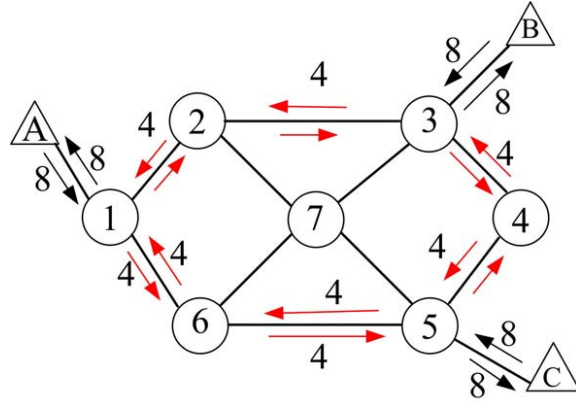

Figure 2: Worst-case scenario of bandwidth reservations.

For each routing, there may be many feasible traffic matrices; in [4] a dual LP was developed as follows.

$$minimize \quad \sum_{i,j} Co_{ij} R_{ij} \tag{1}$$

s.t.

$$\sum_i F_{ik}^{mn} - \sum_i F_{ki}^{mn} = \begin{cases} 1 & \text{if } k = n \\ -1 & \text{if } k = m \\ 0 & \text{if } k \neq m, n \end{cases} \tag{2}$$

$$\forall m, n \in V_c, \forall k \in V$$

$$0 \leq F_{ij}^{mn} \leq 1 \quad \forall l_{ij} \in E, \forall m, n \in V_c \tag{3}$$

$$\sum_{\forall m \in V_c} B_m^- b_{ij}^{m-} + \sum_{\forall n \in V_c} B_n^+ b_{ij}^{n+} \leq R_{ij} \quad \forall i, j \in V \tag{4}$$

5

$$0 \leq R_{ij} \leq C_{ij} \quad \forall i, j \in V \tag{5}$$

$$b_{ij}^{m-} + b_{ij}^{n+} \geq F_{ij}^{mn} \quad \forall m, n \in V_c, \forall i, j \in V \tag{6}$$

$$b_{ij}^{m-}, b_{ij}^{m+} \geq 0 \quad \forall m \in V_c, \forall i, j \in V \tag{7}$$

In this model, R_{ij} is the maximal reserved bandwidth for each link, while F_{ij}^{mn} is the proportion of t_{mn} passing through link l_{ij}. Eq. (2) represents the flow-conservation constraints, which states that all the incoming/outgoing traffic of destination/source is equal to the traffic demand, while there is no net traffic at the intermediate nodes. Two non-negative variables, shadow prices b_{ij}^{m-} and b_{ij}^{m+}, represent the increase in flow on link l_{ij} when increasing B_m^- and B_m^+. Eq. (4) and Eq. (6) ensure that the worst-case traffic matrix can be carried by each link. Eq. (5) is the capacity constraint, which means that a link can not carry more flow than its capacity allows.

4. Top-Down Routing

In RMVP, the high-level topology is composed of the inter-domain links L_{ij}, the CEs, and their corresponding PEs. The intra-domain topology information is known only by the domain itself. Under the hose model, the traffic demand of each CE i is given by the total incoming and outgoing bandwidths, B_i^+ and B_i^-. Based on the traffic demands and limited topology, the objective is to minimize the total cost of the reserved bandwidth for provisioning the L1VPN.

RMVP can be solved by using the LP from Eq. (1) to (7) when the inter-domain and intra-domain topology information is known. Unfortunately, most intra-domain nodes and links are unknown for global optimization. We propose Top-Down Routing (TDR) to solve the RMVP in two steps, as shown in Fig. 3. The first step takes place on the virtual inter-domain topology, which can be generated by knowledge of the inter-domain links and CEs. The second step *refines* the solution, by considering the actual intra-domain topology and utilizing the bandwidth demands that result from the first step. Both steps are solved by using a separate LP, and we propose two approaches to generate the inter-domain topology and perform the routing within each domain.

4.1. Full-Mesh
4.1.1. Inter-Domain Topology
The information we have for routing on the inter-domain topology are the inter-domain links L_{ij}, CEs, and corresponding PEs. As the intra-domain
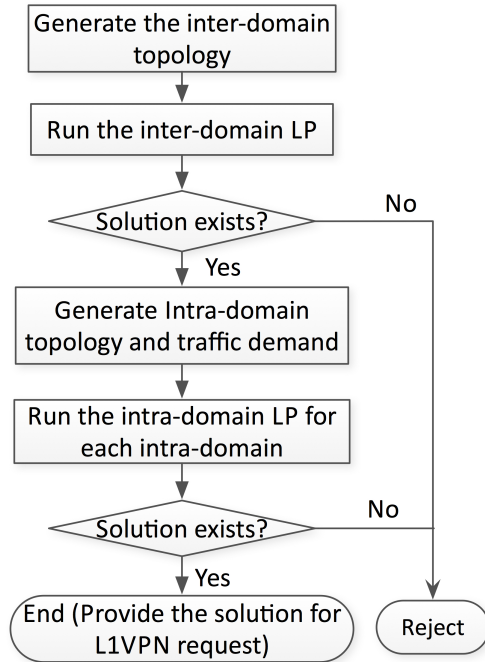
6

Figure 3: Flowchart of TDR routing algorithm.

connectivity is unknown, we generate a full virtual mesh topology in each domain by connecting all known nodes, which are the inter-domain edge nodes and PEs. This is shown in Fig. 4(b), assuming at least one path exists between any two nodes in the same domain.

4.1.2. Routing Strategy

To find the minimum cost for all CEs under the hose model, we first run the LP from Eq. (1) to (7) on the inter-domain topology as top-level routing. Instead of using the single domain set of nodes V, we replace this by the set of inter-domain nodes V_d which contains all the nodes of the inter-domain topology. Based on our objective function, we obtain the minimal reserved bandwidth for each virtual link. The second step is to map the virtual links on paths in each domain V_x. To this end, we use the following LP model in each domain.
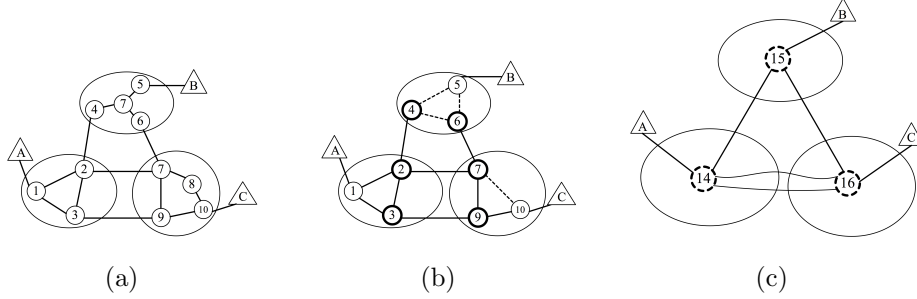
Figure 4: Generating inter-domain topology: (a) original topology, (b) new topology with Full-Mesh, and (c) new topology with Single-Node.

$$minimize \quad \sum_{\forall i,j \in V_x} Co_{ij} r_{ij} \tag{8}$$

s.t.

$$\sum_{\forall i \in V_x} f_{ik}^{mn} - \sum_{\forall i \in V_x} f_{ki}^{mn} = \begin{cases} R_{mn} & \text{if } k = n \\ -R_{mn} & \text{if } k = m \\ 0 & \text{if } k \neq m, n \end{cases} \tag{9}$$

$$\forall m, n, k \in V_x$$

$$r_{ij} = \sum_{\forall m,n \in V_x} f_{ij}^{mn} \quad \forall i, j \in V_x \tag{10}$$

$$0 \leq r_{ij} \leq C_{ij} \quad \forall i, j \in V_x \tag{11}$$

$$f_{ij}^{mn} \geq 0 \quad \forall i, j \in V_x, \forall m, n \in V_x \tag{12}$$

Here, r_{ij} is the bandwidth reserved on intra-domain link l_{ij}. f_{ij}^{mn} is the flow passing through link l_{ij} from demand R_{mn}. Capacity constraints (10) and (11) guarantee that the flow on each link does not exceed the link's capacity. The cost of the virtual links can be calculated by mapping virtual links. Finally, the total cost for the VPN contains the cost of all reserved bandwidth on inter-domain links and intra-domain links.

By running the main LP from Eq. (1) to (7) for the inter-domain topology and full-mesh LP from Eq. (8) to (12) in each domain, we obtain the minimal bandwidth cost based on the limited inter-domain information. The TDR routing strategy will be analyzed by comparing results from the ideal single-domain topology in the next section. The LP used in the second step may be

private and different in each domain, according to the policies or objectives enforced by each InP.

4.2. Single-Node

The technique in Section 4.1 repesents each domain as a full mesh topology. This approach leads to a considerable amount of topology information that needs to be continually distributed and synchronized to ensure accurate path computation results. To reduce this control traffic, we now propose to use a single node to represent each domain, and generate the inter-domain topology by connecting each domain by inter-domain links.

4.2.1. Inter-domain Topology

Even though we can obtain the gateways' information, we simply consider the whole domain as a single node (namely virtual node) in the inter-domain topology. The connections between virtual nodes are inter-domain links, so there may be more than one link between two virtual nodes. As shown in Fig. 4(c), there are two links from node 14 to node 16. We use l_{ij}^{pq} to identify one link in inter-domain topology, where i and j are the domains of the nodes, and p and q are the source and destination nodes of the inter-domain link. For example, links $l_{14,16}^{2,7}$ and $l_{14,16}^{3,9}$ are the two links connecting domains 14 and 16, link $l_{14,14}^{A,1}$ denotes the link from CE A to node 1 in the inter-domain topology.

4.2.2. Routing Strategy

After generating the inter-domain topology with the single-node method, we can calculate the solution with minimal bandwidth for the inter-domain links. To this end, we run the following LP by extending the LP from Eq. (1) to (7).

$$minimize \quad \sum_{\forall l_{ij} \in L_d} Co_{ij}^{pq} R_{ij}^{pq} \tag{13}$$

s.t.

$$\sum_{\forall i \in V_d} \sum_{p,q} F_{ikpq}^{mn} - \sum_{\forall i \in V_d} \sum_{p,q} F_{kipq}^{mn} = \begin{cases} 1 & \text{if } k = n \\ -1 & \text{if } k = m \\ 0 & \text{if } k \neq m,n \end{cases} \tag{14}$$

$$\forall m,n \in V_c, \forall k \in V_d$$

$$0 \leq F_{ijpq}^{mn} \leq 1 \quad \forall i,j \in V_d, \forall m,n \in V_c, \forall p,q \in V \tag{15}$$

9

$$\sum_{\forall m \in V_c} B_m^- b_{ijpq}^{m-} + \sum_{\forall n \in V_c} B_n^+ b_{ijpq}^{n+} \leq R_{ij}^{pq} \quad \forall i,j \in V_d, \forall p,q \in V \tag{16}$$

$$0 \leq R_{ij}^{pq} \leq C_{ij}^{pq} \quad \forall i,j \in V_d, \forall p,q \in V \tag{17}$$

$$b_{ijpq}^{m-} + b_{ijpq}^{n+} \geq F_{ijpq}^{mn} \quad \forall m,n \in V_c, \forall i,j \in V_d, \forall p,q \in V \tag{18}$$

$$b_{ijpq}^{m-}, b_{ijpq}^{m+} \geq 0 \quad \forall m \in V_c, \forall i,j \in V_d, \forall p,q \in V \tag{19}$$

Here, F_{ikpq}^{mn} denotes the proportion of traffic t_{mn} going through link l_{ik}^{pg}, b_{ijpq}^{m-} is the shadow price, which denotes the traffic increasing on link l_{ij}^{pg} when increasing one unit of B_m^-. R_{ij}^{pq} denotes the reserved bandwidth on link l_{ij}^{pg}.

Based on the minimized bandwidth solution for the inter-domain topology, we need to calculate the bandwidth in each domain. The inter-domain link with reserved bandwidth can be considered as a node (namely inter-domain node) with incoming and outgoing bandwidth demands for the respective domain. The set of nodes in new each domain topology added with inter-domain nodes is V_x', while the set of nodes in each original domain is V_x. Consequently, we can compute the intra-domain routing by using the new hose-model traffic demand, coming from the inter-domain solution. To achieve this, we utilize the following LP.

$$minimize \quad \sum_{\forall x \in D} \sum_{\forall i,j \in V_x} Co_{st} r_{st} \tag{20}$$

s.t.

$$\sum_{\forall i \in V_x'} f_{ik}^{mn} - \sum_{\forall i \in V_x'} f_{ki}^{mn} = \begin{cases} 1 & \text{if } k = m \\ -1 & \text{if } k = n \\ 0 & \text{if } k \neq m, n \end{cases} \tag{21}$$
$$\forall k \in V_x', \forall m,n \in V_x' - V_x, \forall x \in D.$$

$$0 \leq f_{ij}^{mn} \leq 1 \quad \forall i,j \in V_x', \forall m,n \in V_x' - V_x \tag{22}$$

$$\sum_{\forall m \in V_x' - V_x} B_m^- b_{st}^{m-} + \sum_{\forall n \in V_x' - V_x} B_n^+ b_{st}^{n+} \leq r_{st} \quad \forall s,t \in V_x', \forall x \in D \tag{23}$$

$$0 \leq r_{st} \leq C_{st} \quad \forall s,t \in V_x' \tag{24}$$

$$b_{st}^{m-} + b_{st}^{n+} \geq f_{st}^{mn} \quad \forall m,n \in V_x' - V_x, \forall s,t \in V_x', \forall x \in D \tag{25}$$

$$b_{st}^{m-}, b_{st}^{n+} \geq 0 \quad \forall m, n \in V_x' - V_x, \forall s, t \in V_x' \tag{26}$$

Here, r_{st} denotes the reserved bandwidth of link l_{st}, and Co_{st} is the cost of bandwidth for link l_{st}. In Eq. (21), f_{ij}^{mn} is the proportion of traffic going through link l_{ij} from m to n. This LP is largely similar to the inter-domain LP, although the traffic demand of the intra-domain LP is obtained from the inter-domain solution.

4.3. Complexity Analysis

We have three routing strategies to get the final solution for L1VPN request, i.e., Single-Domain, Full-Mesh, and Single-Node. The computation complexity of all three models is shown in Table 1. We consider the upper bound of complexity, such that the maximum number of directed links in any topology, consisting of N nodes, is $N(N-1)$. N_d is the number of nodes in the inter-domain topology. N_i is the amount of inter-domain links. After comparison with the single-domain solution, it becomes apparent that TDR introduces additional variables and constraints in each intra-domain model, but reduces the complexity of routing in the inter-domain topology. The single-node TDR has less variables and constraints than full-mesh TDR in inter-domain LP, because N_d is always more than N_i, which means the inter-domain topology of single-node TDR is smaller than the inter-domain topology of full-mesh TDR.

Table 1: Comparison of complexity of TDR LP with single-domain LP.

Routing Strategy	Variables
Single-Domain	$N(N-1)((N_c)^2 + N_c + 1)$
Full-Mesh	$N_d(N_d - 1)((N_c)^2 + N_c + 1)+$ $\sum_{\forall x \in D} N_x(N_x - 1)((N_x)^2 - N_x + 1)$
Single-Node	$(N_i + 2N_c)((N_c)^2 + N_c + 1)+$ $\sum_{\forall x \in D}(N_x)(N_x - 1)((N_x)^2 + N_x + 1)$
Routing Strategy	Constraints
Single-Domain	$2N(N-1)(N_c + 1) + N_c(N_c - 1)(2N^2 - N)$
Full-Mesh	$2N_d(N_d - 1)(N_c + 1) + N_c N_d(N_c - 1)(2N_d - 1)$ $+ \sum_{\forall x \in D} N_x(N_x - 1)((N_x)^2 + 2)$
Single-Node	$(N_i + 2N_c)((N_c)^2 - N_c + 2) + N_c(N_c - 1)N_d$ $+ \sum_{\forall x \in D} N_x(N_x - 1)((N_x)^2 + 2)$

5. Illustrative Numerical Examples

In this section, we first study the performance of the TDR strategy through a special case, and then investigate the impact of different demand profiles on the total bandwidth cost.

5.1. Case Study

Consider three domains in a US-wide network (Fig. 5), and the cost of each link is set as 1. There are 5 CEs connected with PEs 3, 6, 14, 24, and 21. For simplicity, let the demands of the 5 CEs be identical, $B_A^- = B_A^+ = B_B^- = B_B^+ = B_C^- = B_C^+ = B_D^- = B_D^+ = B_E^- = B_E^+ = 4$.

In the single-domain scenario, the complete topology is known. The optimal solution for the L1VPN request is shown in Fig. 5 as a single-domain solution, and the reserved bandwidth for each link is 4 in each direction. This leads to a total cost of 104. Given the symmetric demand profile between CEs, the solution is based on tree routing, as in [7].
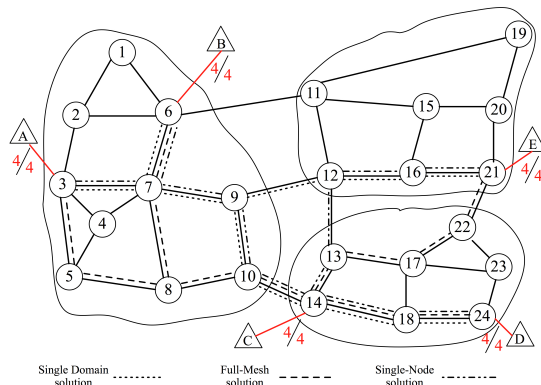

Figure 5: Single-domain and TDR routing.

Using full-mesh of TDR, the optimal solution is shown in Fig. 5 as full-mesh solution. Full-Mesh returns the optimal tree as shown in Fig. 6 as the top-level routing solution, in which the reserved bandwidth of all links is 4. Virtual links $l_{6,10}$ and $l_{10,6}$, as shown in Fig. 6, are mapped on physical links between nodes 6 and 10, which are the optimal mapping for virtual links. However, Full-Mesh is unaware of the existence of nodes 7 and 8 in the top-level topology, thus rendering it impossible to find the optimal tree of the

single-domain solution in Fig. 5. The total L1VPN cost by employing full-mesh strategy is 112, which is only slightly larger than the optimal value. In contrast, the optimal solution of the single-node of TDR approach is shown in Fig. 7. For single-node solution, the bandwidth reserved on the inter-domain links are the same with full-mesh solution which is shown in Fig 6, and the intra-domain costs are also equal with full-mesh solution, so the total cost of single-node solution is 112 as well.
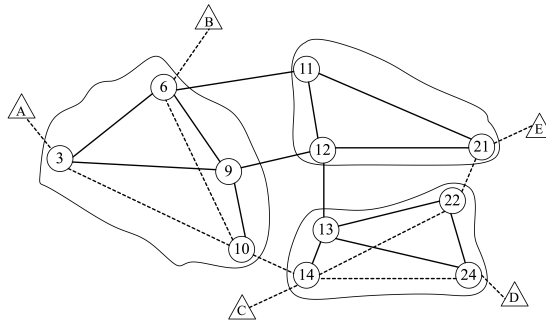

Figure 6: Top-level routing of full-mesh strategy.
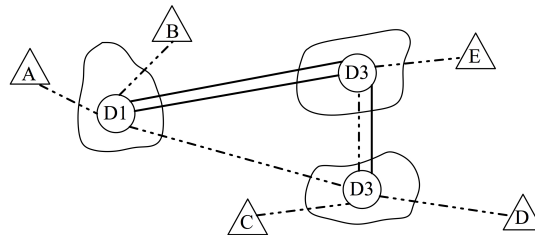
Figure 7: Top-level routing of single-node strategy.

5.2. Global Backbone Network

In the previous section, we found that, in a simple test case, TDR obtains a cost which is close to the optimal value found in the ideal case of single-domain routing. To further demonstrate the features of TDR, we perform more extensive tests on a larger and more realistic scenario. According to [13], we assume that eight locations of a global enterprise request a L1VPN infrastructure for their inter-location networking support. These locations are in

Asia (Beijing and Bengaluru), US (Silicon Valley, Portland, Redmond, and New England), and Europe (Cambridge and Aachen), as shown in Fig. 8. We assume there are 4 domains in this global backbone network which can provide the required bandwidth for the enterprise's L1VPN request, specifically domains US, Europe, India, and China. The respective topologies of these 4 domains are shown in Fig. 9, and the inter-domain links are selected from international connections [14]. The CEs located in Redmond, Portland, Silicon Valley, New England, Cambridge, Aachen, Bengaluru, and Beijing connect with node 1, 2, 3, 19, 27, 32, 63, and 77 respectively. In reality, the costs of
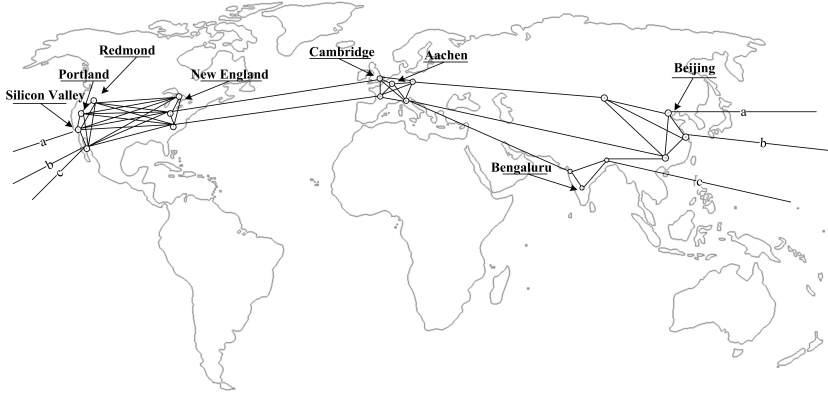

Figure 8: A global enterprise's locations and top-level topology.

inter-domain links are generally higher than intra-domain links. To investigate the features of TDR, we initially set the cost of each intra-domain link as 1, and the cost of each inter-domain link as α. The inter-domain links are generally long-distance connections, so the cost of inter-domain links should be higher than the cost of intra-domain links. We set α to 3 first. The costs of links between CEs and their corresponding PEs are 0, since the reserved bandwidth on these links is identical, irrespective of the deployed routing strategy. Here, we set the capacity of each intra-domain link to 32, and the capacity of each inter-domain link to 96. Below, some scenarios are studied to see how the total bandwidth cost is impacted.

The first experiment compares the total bandwidth cost between the ideal single-domain method and our proposed TDR strategies, for varying symmetric traffic demands. A symmetric traffic demand implies that the incoming and outgoing bandwidth demands for each CE are equal, and, as demon-
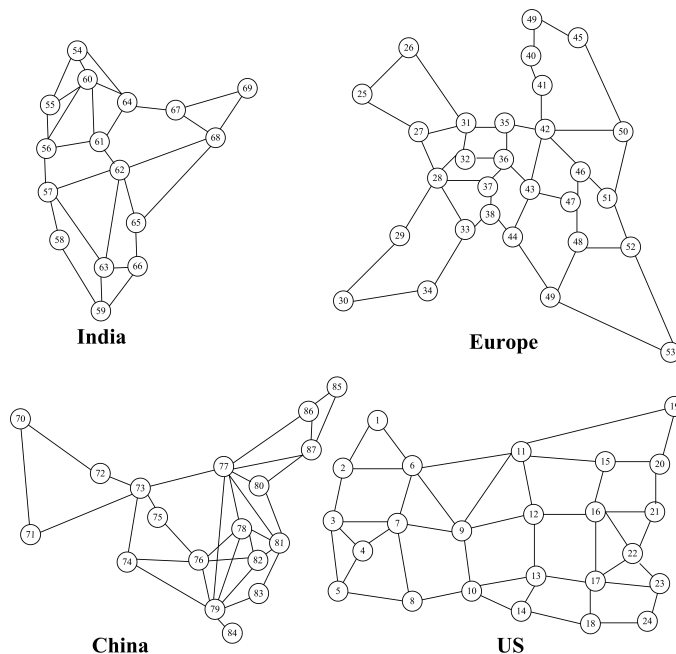
Figure 9: Intra-domain topologies of the considered domains.

strated in [7], this will always lead to a solution based on tree routing when the network capacity is sufficiently high. As shown in Fig. 10, the total cost of all three routing strategies increases along with increasing traffic demand in a linear way. Furthermore, the cost of full-mesh and single-node TDR are slightly higher than the cost of the single-domain approach. Finally, as shown in Figs. 11 and 12, we also calculate the inter-domain and intra-domain cost for each solution separately. The inter-domain cost of full-mesh and single-node are identical, and the inter-domain cost of the single-domain solution is the same as the full-mesh and single-node approaches when the load is below 20, which means 20 units of incoming and outgoing bandwidth for each CE. When the load is 24 or higher for each CE, there is no feasible solution by full-mesh TDR because of the capacity constraint of links. This happens again for single-node TDR when the load is higher than 32 for each CE. The inter-domain cost of single-domain is higher than single-node TDR when the load is higher than 24 because the single-node TDR minimizes the inter-domain cost first and then minimizes the intra-domain cost based on the inter-domain solution, while single-domain minimizes the global cost in a single LP without classifying the inter-domain and intra-domain cost. It
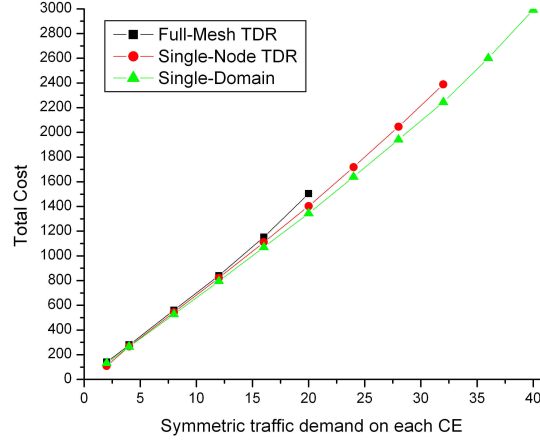
15

Figure 10: Total cost of symmetric traffic demands.

can be observed from Fig. 12 that the intra-domain cost forms the largest contribution to the total cost.
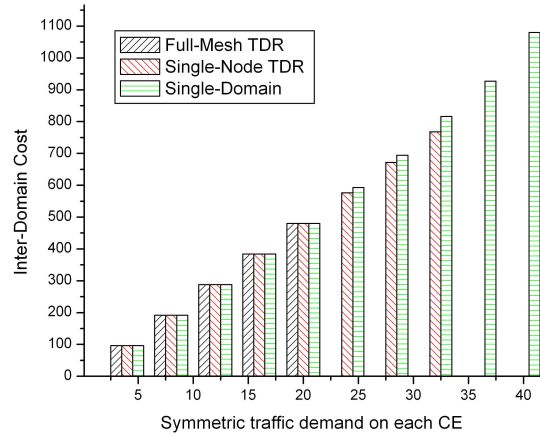

Figure 11: Inter-domain cost of symmetric traffic demands.

We study the impact of different α in the next experiment. We get the cost of L1VPN with different α, and the load of each CE is 4. As shown in Table 2, the cost of inter-domain become 32 with $\alpha = 1$, while the cost is 96 with $\alpha = 3$. However, the cost of intra-domain is the same. We conclude that the topology of L1VPN does not change with different α. Furthermore,
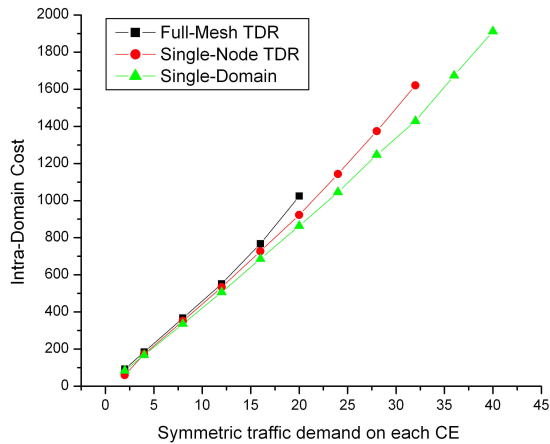
Figure 12: Intra-domain cost of symmetric traffic demands.

Table 2: Comparison of cost with different α.

	$\alpha = 1$	$\alpha = 3$
Full-Mesh Inter-domain	32	96
Full-Mesh Intra-domain	176	176
Single-Node Inter-domain	32	96
Single-Node Intra-domain	172	172
Single-Domain Inter-domain	32	96
Single-Domain Intra-domain	168	168

we study the impact of different amount of inter-domain links. We set α to 1, and set the load of each CE to 4. We cut some inter-domain links from Fig. 8, and get the cost with small number of inter-domain links. The cost in the topology with a large number of inter-domain links is studied as well. As shown in Table 3, the total cost of single-domain becomes less when the number of inter-domain links becomes larger, because there are more feasible solutions for a L1VPN with singe-domain strategy. However, the costs of full-mesh and single-node do not show any pattern, because full-mesh and single-node are two-step strategies which can not consider the intra-domain cost when solving the inter-domain step.

The second experiment compares the cost for symmetric and asymmetric L1VPN demands, implying $B_m^- \neq B_m^+$, while the total incoming and outgoing

17

Table 3: Comparison of cost with different amount of inter-domain links.

	Small	Normal	Large
Full-Mesh Total	280	208	224
Full-Mesh Inter-domain	32	32	32
Full-Mesh Intra-domain	248	176	192
Single-Node Total	208	204	216
Single-Node Inter-domain	32	32	32
Single-Node Intra-domain	176	172	184
Single-Domain Total	208	200	192
Single-Domain Inter-domain	32	32	32
Single-Domain Intra-domain	176	168	152

bandwidth remains constant, i.e., $B_m^- + B_m^+ = B, \forall m \in V_c$. We fix the total bandwidth to B = 16, such that, for example, the first result (left end) in Fig. 13 is tested with $B_m^- = 2$ and $B_m^+ = 14$ for the first group (contains a selection of 4 CEs), and $B_m^+ = 2$ and $B_m^- = 14$ for the second group (contains the remaining 4 CEs). In each experiment, group members are randomly selected; the resulting average over all 5 experiments is shown in Fig. 13. From this result, we observe that, similar to the results for symmetric demands, the inter-domain cost of TDR is the same as the cost of the single-domain approach, and the intra-domain cost of TDR is higher than the intra-domain cost of single-domain approach. For full-mesh and single-domain approaches, the intra-domain cost is unchanged when changing the traffic load. However, the intra-domain cost of single-node is variable along traffic profile.

Figures 14 and 15 detail the reserved bandwidth in the European domain for full-mesh TDR when the traffic demand is $B_{92}^- = 2$, $B_{92}^+ = 14$, $B_{93}^- = 14$, and $B_{93}^+ = 2$. We observe that the inter-domain link cost for full-mesh TDR is unchanged when changing the traffic profile; all incoming and outgoing bandwidth equal 16, irrespective of the traffic profile of nodes 92 and 93. Consequently, the intra-domain cost is the same as well. However, the incoming and outgoing bandwidths of single-node TDR for the European domain are variable and depend on the traffic profile of nodes 92 and 93, since the different inter-domain solutions do not utilize the same link between the European and the US domains. All links between the two domains are equal when the inter-domain LP schedules the flow, so the bandwidth of

two inter-domain links between US and Europe can be assigned randomly. Consequently, the intra-domain cost of the single-node solution is variable and changes with the traffic profile.
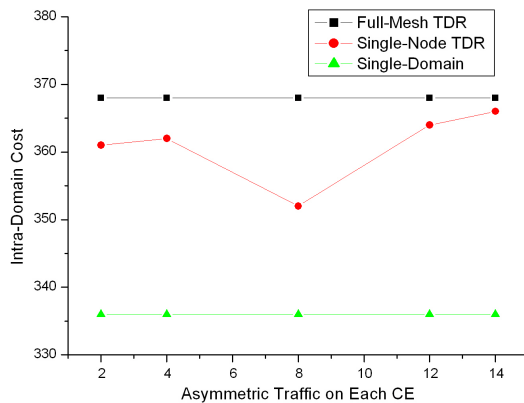

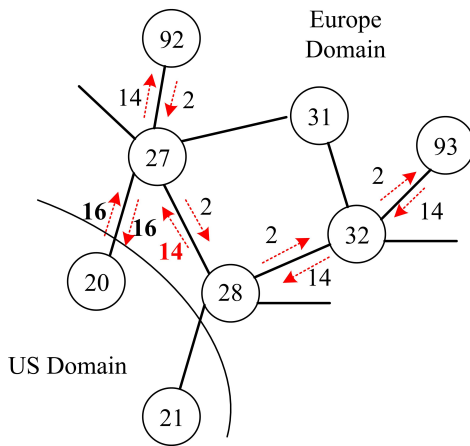Figure 13: Intra-domain cost of asymmetric traffic.

Figure 14: Solution of Europe domain for asymmetric traffic with full-mesh TDR.

6. Conclusion

High-performance routing of L1VPNs in multi-domain backbone networks is a key element to support enterprises for their communication needs. The
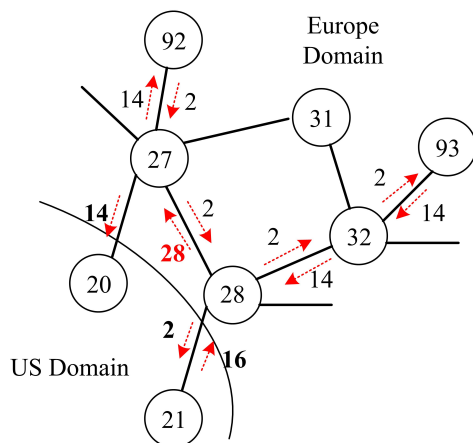
19

Figure 15: Solution of Europe domain for asymmetric traffic with single-node TDR.

high-performance Routing for Multi-domain VPN Provisioning (RMVP) is a fundamental problem for L1VPNs under the hose model. Its main challenges are formed by the limited topology information that is available, and the variability of the traffic matrix. We formulated a model for RMVP utilizing a Mixed Integer Linear Program, and proposed a Top-Down Routing (TDR) strategy to solve the RMVP problem in two steps. Numerical results of a realistic global L1VPN were presented to verify the performance of TDR. By comparing TDR to the ideal single-domain method, we find that TDR is an efficient routing strategy to solve the RMVP problem.

Acknowledgement

References

[1] T. Takeda, I. Inoue, R. Aubin, and M. Carugi, "Layer 1 virtual private networks: service concepts, architecture requirements, and related ad-

vances in standardization," *IEEE Communications Magazine*, vol. 42, no. 6, pp. 132–138, June 2004.

[2] M. Chowdhury and R. Boutaba, "Network virtualization: state of the art and research challenges," *IEEE Communications Magazine*, vol. 8, no. 7, pp. 20–26, July 2009.

[3] N. G. Duffield, P. Goyal, A. Greenberg, P. Mishra, K. K. Ramakrishnan, and J. E. van der Merive, "A flexible model for resource management in virtual private networks," *SIGCOMM Computer Communication Review*, vol. 29, no. 4, pp. 95–108, Aug. 1999.

[4] M. Kodialam, T. V. Lakshman, and S. Sengupta, "Maximum throughput routing of traffic in the hose model," in *Proc. IEEE INFOCOM*, April 2006.

[5] Q. Liu, M. A. Kok, N. Ghani, V. M. Muthalaly, and M. Wang, "Hierarchical inter-domain routing in optical DWDM networks," in *Proc. IEEE INFOCOM*, April 2006.

[6] D. Banerjee and B. Mukherjee, "Wavelength-routed optical networks: linear formulation, resource budgeting tradeoffs, and a reconfiguration study," *IEEE/ACM Transactions on Networking*, vol. 8, no. 5, pp. 598 –607, October 2000.

[7] A. Kumar, R. Rastogi, A. Silberschatz, and B. Yener, "Algorithms for provisioning virtual private networks in the hose model," *IEEE/ACM Transactions on Networking*, vol. 10, pp. 565–578, August 2002.

[8] A. Gupta, J. Kleinberg, A. Kumar, R. Rastogi, and B. Yener, "Provisioning a virtual private network: a network design problem for multicommodity flow," in *Proc. ACM Symposium on Theory of computing*, STOC 2001, 2001.

[9] A. Jüttner, I. Szabó, and R. Szentesi, "On bandwidth efficiency of the hose resource management model in virtual private networks," in *Proc. IEEE INFOCOM*.

[10] T. Erlebach and M. Rüegg, "Optimal bandwidth reservation in hose-model /vpns with multi-path routing," in *Proc. IEEE INFOCOM*, vol. 4, March 2004.

[11] M. Chowdhury, M. R. Rahman, and R. Boutaba, "Virtual network embedding with coordinated node and link mapping," in *Proc. IEEE INFOCOM 2009.*

[12] M. Chowdhury, F. Samuel, and R. Boutaba, "PolyViNE: policy-based virtual network embedding across multiple domains," in *Proc. ACM SIGCOMM Workshop on Virtualized Infrastructure Systems and Architectures*, (VISA 2010), pp. 49–56, 2010.

[13] "Microsoft Worldwide Labs." http://research.microsoft.com/en-us/labs/default.aspx, 2010.

[14] "Submarine Cable Map." http://www.telegeography.com/product-info/map_cable/index.php, 2010.