

A Dynamic Managed VPN Service: Architecture And Algorithms

Ravi S.Ravindran¹, Changcheng Huang¹, K.Thulasiraman²
(1.Carleton University, Ottawa, 2.University of Oklahoma, Norman)

Abstract—VPN as a managed service enables the service provider to offer more demanding and revenue generating services. Some of the more common managed VPN services known today include auto discovery, security and also a potential ability to perform on demand signaling. In this paper we try to tackle an important problem of service providers providing bandwidth service on demand on an IP/MPLS core network. We propose a managed VPN architecture for such a service highlighting the novelty in our architecture. We concentrate on an important aspect of service definition called the topology abstraction service and define a new problem called the VPN core capacity sharing problem that arises in this context. We propose three schemes to solve this problem borrowing established results from graph theory. As part of our simulation study, we evaluate each of these strategies with different call arrival scenarios and present their results.

Keywords: Managed VPN Service, Topology Abstraction

I. INTRODUCTION

With the increasing convergence of transport over routed IP networks, IP-VPN[2] as a managed service has gained a lot of momentum in the last few years. This has mostly been driven by solutions that allow co-existence of L2/L3 VPN service over a single IP/MPLS network. Traditional way to offer IP-VPN services used tunnel based overlay techniques. These VPN's were managed completely by the VPN customers themselves, hence were referred to as Customer Edge (CE) based VPN's. These VPN's employed protocols such as L2TP and IPsec to create tunnels over the L2 or L3 transport layer, over which the VPN payloads were carried. The service provider in this case only provided the capacity and QoS guarantees as agreed upon as part of the SLA, and is ignorant about any VPN existence. A disadvantage often noted with CE-VPN is the lack of scalability with respect to the number of tunnels that had to be maintained at the CE nodes to enable a fully meshed connectivity. The new generation IP-VPN service also called peer based or point-to-cloud based VPN service, requires the Customer Edge (CE) devices to only peer with one or more neighboring Provider Edge (PE) devices. [1] Gives a comparison in terms of performance and scalability of CE based VPN's using IPsec, PPTP, L2TP and peer based approach. The paper also experimentally verifies the superiority of a managed VPN solution in terms of reliability, performance and security characteristics. More recently, IETF

activities involved defining services over MPLS. IP-VPN service definition over MPLS is one of them. [2] Gives a good summary of the managed L2/L3VPN standards work in IETF in recent years. [3] Defines an IP-VPN managed architecture over a core MPLS network. It proposes solutions to two key VPN issues: protocol extensions to differentiate between IP flows belonging to different VPN's, and methodology to route between geographically dispersed VPN sites. To scale the solutions in regards to management complexity, all the VPN extensions are relevant only on the LER's. The core routers or the LSR's are ignorant about any VPN existence.

With the increasing convergence of edge services onto a single IP/MPLS network, future network services will have to address the needs of the potential bandwidth intensive applications of the future, a good example of this is grid applications, where distributed processors might want capacity for a window of time to exchange huge data. In this paper we address one such issue, namely, providing bandwidth on demand dynamic managed VPN service. Before we get into the details of our idea, we briefly review literature on works related to dynamic VPN architectures and SLA schemes differentiated on granularity of traffic demand specifications. The PIPE model [4] SLA tries to emulate layer 2 circuits in a IP-VPN context. Here the virtual link is a pipe connecting two end points of the VPN network. The PIPE is a fixed capacity virtual circuit, and thus bandwidth is committed at any point of the time. This model is the one more prevalent today, and mostly built over L2 transport technologies like ATM ,FrameRelay. This model has two key shortcomings. As the number of VPN endpoints increases, managing the VPN traffic gets complicated, and since each of the pipes is a dedicated circuit, there can be no multiplexing gain among various pipes supporting the VPN or among virtual circuits between VPN's. This scheme requires the VPN's to provide traffic matrix demands. The pipes that traverse the VSP links and nodes can be decided by solving a modified version of Multi-Commodity Flow formulation like in [5] which proposes a multi objective formulation with objective of optimizing resource and link utilization. HOSE model was first proposed in [6]. It tries to alleviate some of the shortcomings of the PIPE model. In the case of a HOSE model the VPN customer specifies a set of end points to be connected with common endpoint with end-to-end performance guarantee. [7] Discusses algorithms for provisioning with

aggregate demand vectors in a HOSE model. RANGE model [8] allows the VPN to specify the requirement as a range of quantitative service, hence the VPN is not required to predict and specify any peak traffic requirements which are typically difficult to anticipate in bursty traffic conditions. [9] Proposes a dynamic programmable VPN architecture that allows spawning dynamic VPN networks with dedicated router and link resources at the discretion of the VPN customer. These works assume access to physical router and link resources through open programmable interface, which may not be possible where strict trust issues exist as in case of an enterprise and an IP-VPN service provider. Also this kind of partitioning requires prior knowledge of demand matrix of the VPN. [10] Proposes dynamic schemes to share resources dedicated to various VPN's statistically taking into account the unused capacity on each virtual circuit and sharing the capacity among newer session arrivals. In all the above schemes the VSP is required to have partial or complete knowledge of the VPN demands, and the capacity is pre-provisioned based on the agreed SLA's. In a dynamic scenario we may encounter two cases, one where the VSP has no ability to predict their future demands, and the other where the future demands are partially known. The other extreme case of full knowledge of traffic demands would be the same as static provisioning and the optimal solutions follow modifications to the multi-commodity flow problem as in [5]. In this study we assume a lack of a demand matrix a priori, and we assume the presence of a set of VPN customers for which the VSP does not pre-provision any resource before hand, but the provisioning takes place on demand. VPN service of this kind, if offered by the VSP, would co-exist with the traditional guaranteed VPN service. In [11][12] we explored the possibility of providing dynamic VPN service where the VSP would provide the VPN's a set of abstract topologies to choose from. These abstract topologies provide a view of the properties of the core network, which is used by the VPN clients in their path computation to determine end-to-end QoS paths. In this paper we propose this concept under the framework of a generalized dynamic managed shared VPN service. We would also like to point out that the concepts discussed in the paper, though centered around L3/L2 VPN's, are equally applicable to L1VPN's as was discussed in [12].

To summarize the rest of the paper, section 2 proposes the idea of dynamic managed shared VPN service, discussing key elements that constitute this service. Section 3 proposes one of the problems called the VPN capacity sharing problem. Section 4 proposes three different algorithms to solve the capacity sharing problem. Section 5 discusses our simulation scenarios and compares the three schemes proposed in section 4.

II. DYNAMIC SHARED VPN SERVICE

In this section we discuss key architectural ideas behind realizing a dynamic shared VPN service. The framework adopted for our architecture is an extension of discussion in [3]. A vanilla set up of a VPN today requires the VPN to

provide the VSP with some traffic statistics, which would be mapped to the core resources permanently. Dynamic VPN service is envisioned for clients that may not require permanent virtual circuits, but require bandwidth that would be of shorter life span. We begin our argument by assuming that the VSP has some portion of resource that it could share among VPN customers seeking bandwidth on demand.

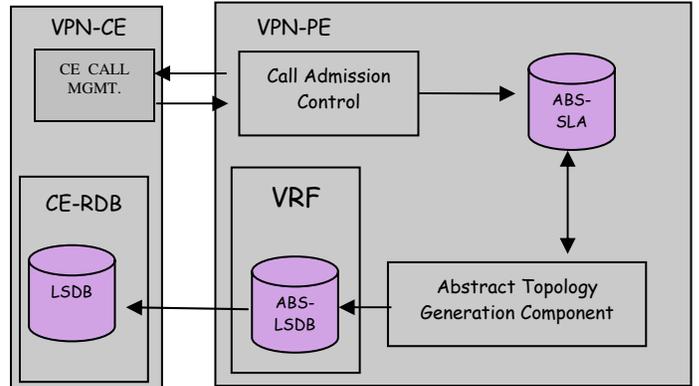


Fig 2.1, Dynamic VPN Service Components

Fig 2.1 shows the components of our architecture. The novel components of our proposal are, the *VPN abstract topology SLA database*, the *abstract topology generation component* and the *VPN abstract topology database* which is specific to a particular virtual routing and forwarding (VRF) instance. The abstract topology components in general deals with the management of the abstract topologies exchanged between the VSP and the VPN. The abstract SLA database is the repository for all the VPN SLA, as defined in [11]. The abstract topology database stores VPN abstractions that are computed by the abstract topology generation component. The key idea behind providing topology abstraction is for the VPN to make a conscious decision before making the bandwidth request. These abstract graphs stored as part of the VRF extensions are part of the special abstract topology database, are flooded by the border nodes to the VPN client nodes, which is then injected into its traffic engineering routing database, shown as CE-RDB in Fig 2.1. When the VPN client router needs to compute a route, it would apply a route computing procedure to decide if a feasible route with the required QoS exists. The route request is then sent to the border node, where another step of call admission control is performed before the call is signaled, the CAC performs the functionality of managing the VPN call statistics, computing end to end paths, and also checking misbehaving VPN'S. In [11] we introduced the notion of an abstraction SLA, and studied various abstract topologies that could be generated as part of the topology abstraction service. In this research we extend the work and use abstraction to virtually partition the core resource, without mapping the shared resource to the link or the switch resource until the call signals the resources along the core nodes and links. Here the capacity shared or dedicated to a VPN is only stored as a state information in border edge nodes. This allows the flexibility to overwrite it at any moment of time, particularly in situations where a VPN needs to get priority

over the resource. The process of providing dynamic abstraction involves several steps, which have been shown in the sequence of events in Fig. 2.6. This begins with the VPN client and VSP agreeing on an abstraction SLA as discussed in [11]. The key components of the SLA would include the type of abstract topology, as an example, for the VSP topology shown serving two VPN's in Fig 2.2, Fig 2.3-2.5 are few abstract topology instances generated for one of the VPN, the metric information associated with the abstract topology, and the abstract topology refresh interval. In [11] we had evaluated the performance of various abstract topologies and a way to provide service differentiation based on the granularity of the abstraction. In this paper we concentrate on the schemes that would be used to share the core capacity between the VPN's, which is a step prior to generating the abstract topologies. The metric we aim to abstract is the VSP's shareable core resource. Once the SLA is agreed upon, the VSP would use one of the proposed abstraction schemes to generate the abstract topology. The remaining steps of Fig.2.6 are self-explanatory. Fig. 2.7 shows the five steps involved in providing a shared dynamic VPN service. These steps are from the point of view of a border node providing abstract topologies to the client. Here, for a given instance of a VPN being serviced by the border node, Steps 1&2 identify the set of border nodes hosting this VPN. In Steps 3&4, we use an abstraction scheme to generate the abstract topology, which is flooded to the respective VPN client served by the border node. This process is repeated every VPN abstract topology refresh interval, which as we stated before, is another abstraction parameter negotiated between a VPN and the VSP.

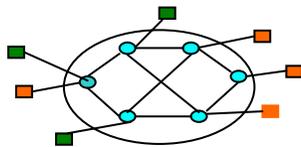


Fig. 2.2: VSP Network

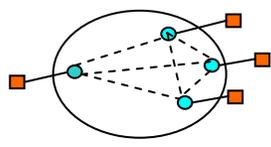


Fig. 2.3: Full Mesh Abstraction

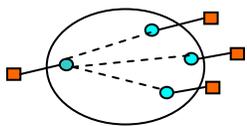


Fig. 2.4: Source Star Abstraction

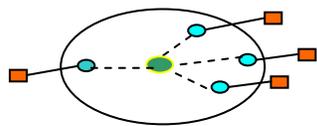


Fig. 2.5: Star Abstraction

The complexity of generating a virtual topology with residual bandwidth as the metric information depends on two key factors: the scheme used to virtualize the core resource and the abstract topology generated out of the virtualization. In [11] we concentrated on the latter part of generating the virtual topologies using the shortest widest path algorithm, which basically generated a shortest widest path tree, which was used to generate abstraction to the VPN clients. But this approach we observed to be too aggressive which resulted in higher crankback calls, which motivated us into research for better heuristics. In the next section we formally define our problem and propose algorithms in the following section.

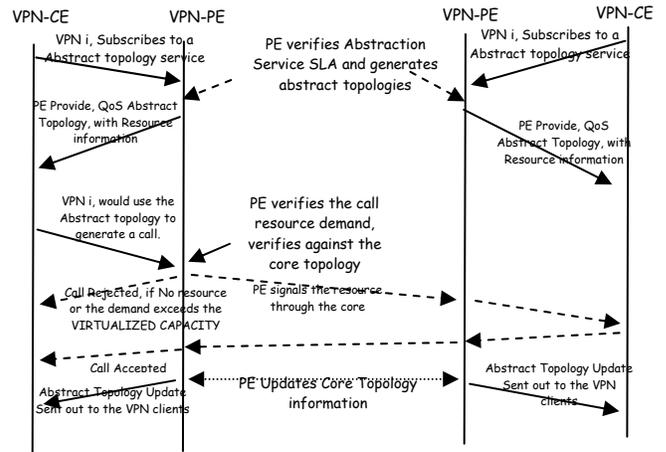


Fig 2.6, Abstract topology Generation Pseudo Code

```

Abstract_Topology( $G, B, VPN_{abs}$ )
 $G$ : VSP Core Graph
 $B$ : Set of border node
 $V_{abs}$ : Set of subscribed VPN's
Step 1: For each VPN  $k \in VPN_{abs}$ 
Step 2: Find set  $PE_k \subset B$  and  $CE_k \subset C$  for VPN  $k$ 
Step 3: Apply a proposed Abstraction Scheme to
generate Abstract Graph  $G_{abs(k)}(V, E)$ 
Step 4: Update VPN node's  $CE_k$  with  $G_{abs(k)}(V, E)$ 
Step 5: Goto Step 1

```

Fig 2.7, Topology Abstraction Process

III. VPN CAPACITY SHARING PROBLEM

We start by introducing the notations used in the rest of the paper: Given a directed graph $G(V, E)$ representing the core of the VSP, each link $e_{ij} \in E$ is associated with total capacity of l_{ij} , and shared capacity represented as R_{ij} . $R_{ij} - l_{ij}$ represents the resource use for guaranteed service. Let B represent the set of border PE nodes, and C the set of all CE nodes connected to the VSP's network. Let the set VPN_{abs} be the set of customers subscribing to abstraction service. A border node $b_i \in B$ supports one or many of the VPN instances from set VPN_{abs} . The VPN's are connected to the border nodes through their CE nodes. For a VPN_{abs} instance k we represent the set of corresponding CE nodes as set CE_k and the set of border nodes as PE_k , and for a given border node $b_i \in PE_k$ the set of CE nodes connected to it as $CE_{i,k}$. As part of the dynamic VPN service, each of the VPN's in VPN_{abs} is served with abstract topologies. For a given VPN_{abs} i we represent the abstract graph as $G_{abs(i)}(V_i, E_i)$, where V_i is the set of virtual nodes, some or all of which may map to a border node PE_i . E_i is the set of virtual link connecting a pair of virtual nodes $(x, y) \in V_i$. A virtual link $e \in E_i$ is associated with a vector of abstracted metric denoted as $v(e)$. Here we restrict ourselves to only one abstracted metric i.e virtualized core capacity. Hence for edge $e \in E_i$ connecting nodes $(x, y) \in V_i$, we have $v(e) = \{bw_{abs}[x][y]\}$. This bandwidth represents the capacity exposed by the VSP between the pair of V_i nodes connected by a virtual link. The capacity over an access link is the available physical residual capacity. The VPN's use the abstract graph $G_{abs(i)}(V_i, E_i)$ to compute end-to-end path. Using the above notations we state the VPN capacity sharing problem as follows:

Given a set of VPNs in VPN_{abs} , where each VPN instance $i \in VPN_{abs}$ is provided by an abstract topology $G_{abs(i)}(V_i, E_i)$, and each virtual link is assigned a virtual capacity as decided by the VSP. The objective of the problem is to devise a methodology for the VSP to share the VSP core resource among VPN_{abs} instances so that for a given VPN, the VSP maximizes its probability of making a correct decision of successfully computing or rejecting a path locally.

The above problem's objective has been defined from a VPN customer's perspective. It can also be framed from the VSP's perspective where the objective is to share the core capacity so as to maximize the network utilization. In case of a dynamic VPN service maximizing network utilization is not only a function of the type of routing strategy applied in the core network, but also function of efficiency of the scheme applied to share the shareable core capacity. A poor capacity-sharing algorithm might just be very conservative in exposing the core resource information, and hence could lead to bad resource utilization. Generating an abstraction to a VPN is a two-step process. First step involves computing a subgraph $SG_i(V_i, E_i)$ where V_i is the set of PE nodes and a subset of the core nodes, and E_i the subset of physical core links. The second step is to generate an abstract topology $G_{abs(i)}(V_i, E_i)$ from the sub-graph $SG_i(V_i, E_i)$. Since the metric to be abstracted is the residual capacity, the complexity of generating a topology boils down to the complexity involved in generating a subgraph for each of the VPN's. For a simple case where the subgraph is a tree, this problem boils down to generating Steiner Tree graph with the link weights as a function of the residual capacity. This problem is known to be NP-Complete and makes the core capacity sharing problem difficult to solve. In the next section we propose three abstraction heuristics for the capacity sharing problem, which is to be used in Step 3 of the abstraction process described in Fig 2.7. In the rest of the discussion we have assumed $V_i \subseteq PE_i$, hence using the terminology border nodes instead of virtual nodes in the context of abstract topologies in the rest of the paper. In Section 5 we evaluate these heuristics using a simulation setup.

IV. VPN CAPACITY SHARING ALGORITHMS

In this section we propose capacity sharing schemes that use different graph algorithms to provide abstract views of the VSP's core network. The motivation behind the abstraction schemes is to provide accurate information to the VPN customers by minimizing overlap between the virtualized core resource and at the same time maximizing the capacity multiplexing gain in the VSP's core. The key challenge of any capacity sharing schemes is to fairly share the core resource among all the VPN's and to minimize the call rejections at the associated PE nodes. We propose three approaches. The first one uses maximum capacity path for abstraction. This was suggested in [11] too. The two other novel approaches are called the max-min bound approach and tree graph approach. In the following we elaborate on each of these schemes.

Maximum Capacity Abstraction Scheme:

In this approach the VSP exposes the capacity of the widest path between two border nodes belonging to a given VPN. For

a given pair of border nodes $(b_1, b_2) \in B$ belonging to a given VPN i , let $P = \{p_1, p_2, \dots, p_k\}$ be the set of the k paths available between the two nodes. Let $C(p_i)$ be the capacity of each path. We use $\max_{i=1 \dots k} (C(p_i))$ as the bandwidth $bw_{abs(b_1, b_2)}$ between

the two nodes. This problem can be solved in $O(V^2)$ time using an algorithm suggested in [13], using which a shortest widest path tree can be computed from a border node computing the abstraction. Once the tree is computed, any desired abstraction can be generated based on the agreed abstraction SLA [11]. The total complexity of this scheme for a single VPN would be $O(|V|^2 + |ABS|)$, here $|ABS|$ refers to the complexity involved in generating the abstract topology. This approach is very aggressive since the same maximum capacity path is abstracted to all the VPNs. This aggressive mode of capacity sharing does well when the abstract topologies are accurate at the time calls are being processed. There are two reasons why this scheme might not fair well. First is a case where one of the VPN's starts misbehaving and seeks bandwidth more aggressively. Another reason why this may not work well is when multiple calls from different VPN's arrive at the border node at the same time leading to resource contention and to higher call crank back probability. The next heuristic we propose is more conservative, which also addresses the key issue handling multiple call arrivals from different VPN's simultaneously.

Mixed Bound Abstraction Scheme:

In this scheme, for a VPN k the VSP provides two capacity bounds, upper bound capacity ($UBOUND_k$) and lower bound capacity ($LBOUND_k$). We define $UBOUND_k$ as the aggregate bandwidth that can be requested in between two consecutive abstract topology update instances. $UBOUND_k$ also indicates the slice of the shared network capacity available between any two border nodes. This bound is derived from the max-flow computation, which does not map to a single path flow. $LBOUND_k$ on the other hand gives a single path flow capacity approximation that can be requested from the VSP. The $UBOUND_k$ as said is computed using the max flow algorithm. The key idea here is to virtually expose a slice of the max flow possible between two border nodes for a given VPN. The computed max flow is shared with the VPN's as decided by their priority factor α_i for VPN i . For a given border node k , α_i satisfies, $\sum \alpha_i = 1$. Here i sums over all the VPNs supported on the border node. For example, for a given VPN k , and a pair of border nodes $x, y \in PE_k$, if $bw_{max}[x][y]$ is the maximum flow achievable between the two border nodes, then the VSP would expose a capacity of $UBOUND_k[x][y]$ equivalent to $\alpha_i * bw_{max}[x][y]$ to VPN i . The priority factor could be used to control the net capacity virtually allocated to a given VPN at any point of time. By default we assign it to the number of subscribed VPN's, i.e $\alpha_i = 1/|VPN_{abs}|$.

In order to give a single path flow approximation, we compute $LBOUND_k$. In an aggressive mode, the $LBOUND_k[x][y]$ can be set to the capacity of the shortest widest path, but with the drawback of higher crank back calls which are caused by multiple calls arriving simultaneously at the border node with

aggregate demand exceeding the available capacity. Here we propose a less aggressive approach using M-Route flow algorithm to compute the lower bound. We would suggest interested readers to refer [16] for definition of M-Route flow. To define loosely, M-Route flow is any flow that can be expressed as a non-negative linear sum of elementary M-Flows. [16] Defines M-Route flow in the context of both edge and vertex disjoint paths. For our study we use the theory of M-Route edge flows. Using M-Route flow as the LBOUND the border node has a higher probability to satisfy requests simultaneously for aggregate capacity less than or equivalent to M-Route max flow, this hinges on the fact that a M-Route max flow can be decomposed into set of M-Route elementary flows. The question now is a way to choose the value of M. M is upper bounded by the maximum number of edge disjoint paths existing in a given VSP core topology, which can be obtained applying Menger's theorem. Ideal way to assign a value to M would be to analyze the call arriving pattern, and fixing it to the expected number of calls arriving simultaneously at the border node during a time window of the abstract update interval, because its most likely that the VPN's are making requests using the same abstract topologies in this time period. But if the call arrivals are very aggressive then one may find M to be still very large, which may return zero flow values because of physical topological constraints. Another practical way to assign M would be to initialize it to the number of VPN instances being served by the concerned border node. In order to correlate UBOUND_k and LBOUND_k for a given VPN, the amount of LBOUND_k exposed is made a function of α_i , hence if $bw_{m-route}[x][y]$ is the maximum M-Route flow achievable between the two border nodes, then the VSP would expose a capacity of $LBOUND_k[x][y]$ equivalent to $\alpha_i * bw_{m-route}[x][y]$ to VPN i . [16] Proposes an algorithm to compute a M-Route Flow, which can be found in maximum of (M-1) runs of the max-flow algorithm bringing the pseudo polynomial complexity to $O(M*|V|^2|E|)$ using the shortest augmenting path algorithm. For a given pair of border nodes x and y , we can observe that for $M=1$, $UBOUND_k[x][y] \geq LBOUND_k[x][y]$, but this condition may not hold for $M \geq 2$. In cases where $UBOUND_k[x][y] < LBOUND_k[x][y]$, we set $UBOUND_k[x][y] = LBOUND_k[x][y]$. In cases where $LBOUND_k[x][y]$ is zero, we switch to a best effort service and set $LBOUND_k[x][y]$ to the maximum capacity value between x and y . Since the M-Route flow dominates the complexity total complexity for the mixed bound approach is $O(|B|*|M|*|V|^2|E| + |ABS|)$. Fig 4.1 shows the pseudo code to compute the UBOUND_k and LBOUND_k for a given VPN k from a border nodes perspective.

Mixed Bound Abstraction($G, VPN(k), B_b, C_{i,k}, \alpha_k, M$)

Step 1: Find border node set $B_{i,k}$

Step 2: Init. Graph $G_{abs(k)}(B_{i,k}, E_k)$. For $\forall b_j \in B_{i,k}$ $e(B_i, b_j) \in E_k$, $UBOUND_k[B_i][b_j] = 0$, $LBOUND_k[B_i][b_j] = 0$

Step 3: For each virtual link $e_{x,y} \in E_k$, where $x = B_i, \forall b_j \in B_{i,k}$
Set $UBOUND_k[B_i][b_j] = \alpha_k * MaxFlow(G, x, y)$
Set $LBOUND_k[B_i][b_j] = \alpha_k * M-RoutFlow(G, x, y, M)$
If ($LBOUND_k[B_i][b_j] = 0$)
Set $LBOUND_k[B_i][b_j] = MaxCapacity(G, x, y)$
If ($UBOUND_k[B_i][b_j] < LBOUND_k[B_i][b_j]$)
Set $UBOUND_k[B_i][b_j] = LBOUND_k[B_i][b_j]$

Step 4 : Flood $G_{abs(k)}(B_{i,k}, E_k)$ to nodes $C_{i,k}$

Fig 4.1, Mixed Bound Abstraction Scheme

Tree Based Abstraction Scheme

Here we partition the network using tree graphs, in such a way that there is minimum overlap between the trees computed for each VPN used to build abstract topologies. The tree approach guarantees single path abstractions, at the same time giving better guarantee of finding a path when requested for the core capacity, as long as there is a minimal overlap of the VPN trees. In order to minimize the VPN tree overlaps, we make the edge weight as a function of number of occurrences. We explain the abstraction scheme following the steps stated in the pseudo code as shown in Fig. 4.2. In *Step 3*, before computing the Steiner graph, we initialize the link weights as function of the residual capacity. In this abstraction scheme we also introduce a new link variable called *vpnCount*. This variable indicates the number of times the edge has been part of a Steiner tree. We initialize the edge cost as a function of *vpnCount* and the residual capacity. The idea of resetting the cost is that, the links with more available bandwidth would be at lower cost compared to links with high cost and higher utilization. Since the goal is to optimize the usage of the core resource, we employ a minimum cost tree algorithm, which also correlates with preventing over subscribing congested links, leading to even abstraction of the available resource. In *Step 4* we compute the Steiner tree, T_i , for a VPN i and *Step 5* shows the *vpnCount* variable of the edge $e_i \in T_i$ being incremented by 1. Making the link cost a function of *vpnCount* also dissuades the future VPN tree computations from using the links used by the previous VPN tree computation. After the trees are computed, in *Step 6* we assign the virtual capacity to the virtual link of the abstract graph. For a particular VPN i , and a source root node which is also the concerned border node b_i , we summarize bandwidth between the node b_i and another border node b_j , by identifying the edge set E_i from the tree T_i that comprises the tree path. The virtual capacity of the path is the bottleneck capacity of all the edges belonging to the path. Steiner tree is a strongly NP Complete problem, but literature provides good heuristics to compute Steiner trees. We use the minimum cost Steiner tree algorithm from [15] for the tree computation. The complexity of deriving a source rooted abstract topology for a given VPN is $O(|B_k|*|V|^2 + |ABS|)$.

Steiner Tree Abstraction Scheme($G, VPN(k), B_b, C_{i,k}$)

Step 1 Find Border Node set $B_{i,k}$

Step 2 Initialize Graph $G_{abs(k)}(B_{i,k}, E_k)$
For $\forall b_j \in B_{i,k}$, $e(B_i, b_j) \in E_k$, Set $bw_{abs}[B_i][b_j] = 0$

Step 3: For edge $e_i \in G(V, E)$,
Set $Cost(e_i) = vpnCount(e_i) * (I(e_i)) / (R(e_i))$

Step 4: Use a Steiner tree heuristic on nodes B_i to generate Steiner graph $T_i(V, E)$

Step 5: For each edge $e_i \in T_i(V, E)$ and $e_i \subset G(V, E)$
Set $vpnCount(e_i) = vpnCount(e_i) + 1$

Step 6: For each $b_j \in B_{i,k}$, Let $P = \{e_1 \dots e_k\}$, $e_i \in T_i(V, E)$
Set $bw_{abs}[B_i][b_j] = \min\{bw(e_1) \dots bw(e_k)\}$

Step 7: Flood $G_{abs(k)}(B_{i,k}, E_k)$ to nodes $C_{i,k}$

Fig 4.2, Tree Graph Abstraction Scheme

V. SIMULATION STUDY

The simulation was implemented using OpNet, on a network topology of 50 nodes with an average degree of 4. 10% of the total number nodes in the core topology n were chosen randomly as the PE nodes. Each of the PE nodes was configured to handle five different VPN instances. For comparing the capacity sharing schemes we use four different metrics: *Success Ratio*, *CrankBack Ratio*, *MissCall Ratio* and the *Call Performance Ratio*. *Success Ratio* is a measure of making a right routing decision using the abstraction provided by the VSP. This includes the case of computing a feasible path, as well as the case of rejecting a path locally at the VPN client's end. The correctness of a rejected path is verified by re-computing the path with the exact state of the network. The *CrankBack Ratio* is defined as the fraction of calls that have been cranked back because of a successful path computation at the VPN's end, but bouncing back from the border node because of insufficient capacity in the core. Note that though crankback ratio and success ratio are correlated they are not exactly complementary. *MissCall Ratio* is the fraction of calls that have been rejected wrongly locally by the VPN, even when the core has sufficient resource to accept it. This could be either because of inaccurate abstraction of core capacity or poor granular abstract topology subscribed by the VPN client. One can see that success ratio and crankback ratio are correlated, and the ideal values one expects is 1 for the success ratio and 0 for the crankback ratio. An aggressive heuristic may have very good success ratio but a poor crankback record. To bring out this relationship in assessing the performance of the heuristics, we introduce the Call Performance Ratio metric. We define the *Call Performance Ratio* of the heuristic as the ratio of success ratio to the crankback ratio. So ideally, higher this value betters the heuristic. To make an objective study of the differences in these capacity sharing heuristics, for our simulations we fixed the abstract topology to be source-rooted abstraction as described in [14] and also discussed in [11]. The simulation setup has link bandwidths for the access as well as the core initialized to 1000 units. The bandwidth call requests from the VPN client nodes were modeled as Poisson arrivals. Similarly the call holding times are negatively exponentially distributed. During the simulation it was assumed that all the core border nodes have up to date information of the core topology. The VPN abstract topology update interval was set at value less than the mean arrival rate of the calls. The bandwidth request follows a uniform distribution of [1-500]. In the graphs discussed later the x-axis traffic load relates to the ratio of the arrival rates to the service rate.

Fig 5.1 compares the Success ratio for each of the three abstraction algorithms. Here we see the mixed bound scheme performing poorly than the other two schemes, the reason being that a lot of calls were rejected locally because of the aggregate demands exceeding upper bound capacity as dictated by the VSP, even though there is resource available in the core. The maximum capacity scheme and the tree graph scheme of abstraction show a much better Success ratio compared to mixed bound scheme. Though we notice a slightly poorer performance by the tree graph approach, this is offset by a very low crank back ratio that we see later, which makes it in fact a better algorithm. Fig. 5.2 gives a comparison of the CrankBack

Ratio of the three schemes. As expected, we observe that the mixed bound scheme has the least CrankBack ratio compared to the other schemes,. The lower CrankBack ratio can be attributed to the upper bound capacity that can be used to determine the aggregate capacity the VPN could seek from the core. Hence most of the calls are rejected locally when the demands exceed the upper limit capacity. Comparing the mixed bound approach and tree graph approach, we see that the tree graph approach outperforms the mixed bound approach in regards to the CrankBack ratio, which can be attributed to better abstraction strategy applied by abstracting capacity. Fig 5.3 compares the MissCall ratio of the three schemes. We see that in most cases the total number of missed calls are twice the ratio of the mixed and the tree graph schemes. The higher number of missed calls can be attributed to higher calls being rejected locally by the VPN's even though there is the needed resource to satisfy the requests. Fig 5.4 gives a comparison of the performance metric, which gives a clearer idea of the heuristic's performance. We see that of the three heuristics the mixed bound heuristic performs the best, followed by the tree graph heuristic. But the performance of the Mixed Bound heuristic is offset by the complexity of the algorithm which is at least $O(|E|)$ times expensive than the Maximum Capacity or the Tree Graph heuristic.

In addition to comparing the three parameters discussed above, we also evaluated the core network utilization when each of the three schemes is employed. In Fig. 5.5 we note that the maximum capacity abstraction leads to the best network utilization. The higher network utilization for the maximum capacity scheme can be attributed to its aggressive nature of sharing resources, but with a drawback of higher CrankBack ratio. Tree graph scheme does slightly poorly compared to the maximum capacity approach. This also shows that in a core resource sharing scenario the network utilization is also a function of the capacity sharing heuristic. Summarizing the results, we see that of the three schemes, aggressive schemes like maximum capacity heuristic enjoys higher success ratio, but with the drawback of high CrankBack ratio. Mixed bound scheme turns out to be more conservative, leading to poor Success ratio and high miss call ratio, but with very good CrankBack ratio. Tree graph approach enhanced with the idea of tree separation when abstract topologies are computed stands out satisfactorily in all the four performance metrics, which makes it the best of the three schemes.



Fig 5.1, Success Ratio Comparison

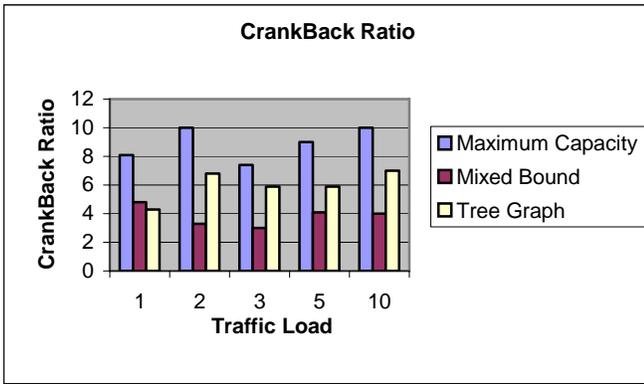


Fig 5.2, CrankBack Ratio

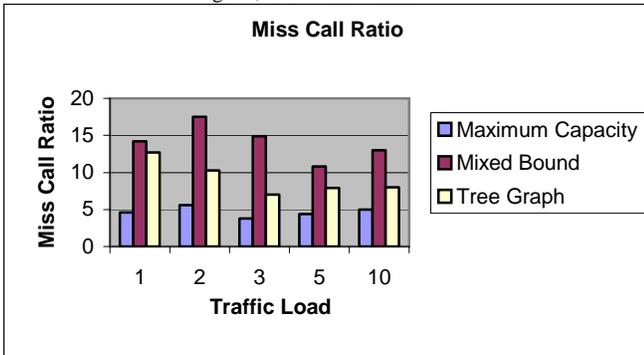


Fig 5.2, Miss Call Ratio

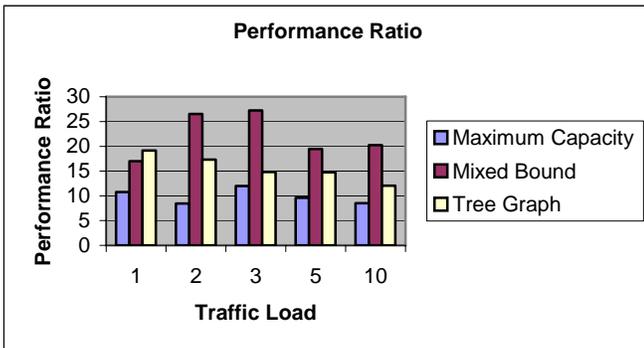


Fig 5.2, Performance Ratio

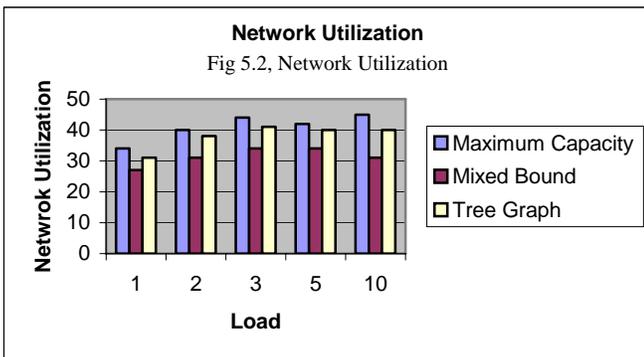


Fig 5.4, Network Utilization Ratio Comparison

VI. CONCLUSIONS

In this paper we proposed a dynamic managed VPN service and noted its differences from the traditional VPN definitions.

An architecture as an extension to the existing IP-VPN solution was proposed. We then defined core capacity sharing problem in a dynamic managed VPN service context. As a way to solve this problem we proposed three heuristics. The maximum capacity abstraction, which is an aggressive way to sharing resource, has satisfying Success ratio, but its CrankBack ratio is almost twice those of the other two schemes. Mixed Bound approach tries to address the drawback of the maximum capacity scheme by defining a virtual upper bound and lower bound for the bandwidth that can be requested from the VSP. This scheme was the most conservative of the three schemes with poor Success Ratio and high algorithm complexity. The third scheme we proposed uses Tree Graphs, using Steiner trees as a way to virtualize capacity applying a method to minimize overlap of the trees that were later used to generate abstract topologies. This method faired well in all the four performance metrics. It also performed quiet well in terms of overall network utilization of the core.

VII. REFERENCES

- [1] Francesco Palmieri, "VPN Scalability Over High Performance Backbones evaluating MPLS VPN against Traditional Approaches", Proceeding ISCC 03
- [2] Paul Knight, Chris Lewis, "Layer 2 and 3 Virtual Private networks: Taxonomy, Technology and Standardization Efforts", IEEE Communication Magazine, June 2004
- [3] IETF draft, Eric Rosen et al, "BGP/MPLS IP VPN's"
- [4] Nick Duffield et al, "draft-duffield-vpn-qos-framework-00.txt", IETF Draft.
- [5] Chun Tung Chou, "Traffic Engineering for MPLS-based Virtual Private Networks", IEEE, 2002
- [6] N.G.Duffield, P.Goyal, A.Greenberg, P.Mishra et al, "A flexible model for resource management in VPN", in Proc. ACM SIGCOMM, 1998, pp 95-108
- [7] Amit Kumar, Rajeev Rastogi, Avi Siberschatz, "Algorithms for Provisioning Virtual Private Networks in Host Model", IEEE/ACM Transactions on Networking, Vol. 10 No.4 August 2002
- [8] Ibrahim Khalil , Torsten Braun, "Edge Provisioning and Fairness in VPN-DiffServ Networks", IEEE-ICC 2000
- [9] Rebecca Issacs, " Light Weight Dynamic Programmable VPN", IEEE-OPENARCH, 2000.
- [10] Rahul Garg, Huzur Saran, "Fair Bandwidth Sharing about Virtual Networks: A capacity Resizing Approach", IEEE, INFOCOM 2000
- [11] Ravi Ravindran, Peter Ashwood-Smith et al, "Multiple Abstraction Schemes for Generalized Virtual Private Networks", IEEE-CCECE-2004, Niagara.
- [12] Ravi Ravindran, ChangCheng Huang, K.Thulasiraman, "Topology Abstraction as VPN Service", IEEE, ICC, 2005.
- [13] Z.Wang and J.Crowcroft, "Quality-of-Service Routing for Supporting Multimedia Applications," *IEEE Journal on Selected Areas in Communications*, Vol. 14, no.7, September 1996, pp.1228-1234.
- [14] Turgay korkmaz and Marwan Krunz, "Source-Oriented Topology Aggregation with Multiple QoS parameters in Hierarchical Networks", ACM Transactions on Modeling and Computer Simulations, Vol 10, No. 4, Pages 295-325, October 2000.
- [15] Kou, L., G. Markowsky, L.Berman, "A fast Algorithm for Steiner Trees", Acta Informatica, Springer-Verlag, 1981: vol. 15, pp.141-145.
- [16] W. Kishimoto, M.Takeuchi, "On M-Route Flow in Networks", ICCS/ISITA'92.