

Bandwidth Requirements for Protected VPNs in the Hose Model

Ashok Balasubramanian and Galen Sasaki¹

Department of Electrical Engineering

University of Hawaii at Manoa

Honolulu, HI 96822, USA

{ashok, sasaki}@spectra.eng.hawaii.edu

Abstract — A Virtual Private Network (VPN) is a service by a telecommunication or IP service provider to emulate a private network for a customer. In the *Hose Model* [1], the customer specifies only the maximum traffic rate that can originate/terminate at each VPN site instead of the complete traffic matrix. Protection of hose VPNs from single link failures is considered. Simulation results are presented that compare the bandwidth of line and path protection.

I. INTRODUCTION

A VPN is a service provided by a telecommunication or IP service provider (SP) to emulate a private network, that connects a collection of sites. In the *Hose Model* [1], each VPN site k is connected to the SP through a single port with capacity b_k . The VPN can support all traffic matrices such that the rate of aggregate traffic originating or terminating at each site k is at most b_k . In [2], the VPNs were restricted to tree topologies, and a polynomial time algorithm was presented that computes trees with minimum total link bandwidth. In [3], protecting such VPNs from single link faults was considered. Computing the minimum protection bandwidth was shown to be NP-Complete. A polynomial-time approximation algorithm was given but the bounds in the worst case are 16 times larger than optimal.

We also consider protecting a hose VPN tree from single link failures. Its SP network has a topology G with bidirectional links, that have essentially infinite capacity. The sites of the VPN are located at a subset of nodes W of G . Each site $k \in W$ has at most b_k originating/terminating traffic. The topology of the VPN is a tree T , where W are its leaf nodes. It is assumed that the VPN tree is computed by the optimal algorithm in [1]. The bandwidth for a link $(u, v) \in T$ is determined as follows. If we delete (u, v) from T then we have two subtrees $T(u, v)$ and $T(v, u)$, where the first is rooted at u and the second is rooted at v . Let $B_{uv} = \sum_{k \in T(u, v) \cap W} b_k$, which is the maximum total traffic originating (and terminating) from $T(u, v)$. Similarly, we can define $B_{vu} = \sum_{k \in T(v, u) \cap W} b_k$ for $T(v, u)$. The provisioned bandwidth for the link $(u, v) \in T$ is $c_{uv} = \min(B_{uv}, B_{vu})$. In the next section, we consider *line* and *path protection*, where the *protection bandwidth* (i.e., spare bandwidth to restore the network in case of failure) is shared.

II. LINE AND PATH PROTECTION

For line protection, each link $(u, v) \in T$ has a pre-computed back-up path p , that connects u to v while avoiding link (u, v) . Since p is a replacement for (u, v) , protection bandwidth of c_{uv} must be available on all its links.

For path protection, each possible failure $(u, v) \in T$ has a pre-computed back-up path p' that reconnects the surviving subtrees $T(u, v)$ and $T(v, u)$. Protection bandwidth must be allocated along a path p , that starts from u , follows $T(u, v)$ to p' , crosses p' , and then follows $T(v, u)$ to v . The protection bandwidth along p' is $c_{u, v}$. The protection bandwidth for a link (x, y) in $T \cap p$ is for traffic that normally goes through (u, v) but is rerouted through p' when the link fails. For a link $(x, y) \in T(u, v) \cap p$, where u is closer to x than y , the protection bandwidth is $\min(B_{xy} - B_{vu}, B_{yx} + B_{vu}) - c_{xy}$. The protection bandwidth of links in $T(v, u) \cap p$ can be derived similarly.

Figure 1 shows simulation results when the SP topology G is the NSFNET-16 topology (16 nodes, 25 links). The results are for different numbers of VPN sites for line protection (LINE), path protection (PATH), the optimal hose VPN tree without protection bandwidth [2](TREE), and the approximation algorithm of [3] (APPR). To determine the bandwidth for line and path protection, mixed integer linear programming problems were formulated and solved. Each data point is the average total bandwidth of 10 instances, where the VPN sites and b_k values were chosen randomly, and the b_k values were chosen uniformly in the interval $[2, 100]$.




Fig. 1: NSFNET-16 topology and average VPN bandwidths.

III. CONCLUSIONS

Path protection has much lower bandwidth than either line protection or the approximation algorithm in [3]. In the case of the NSFNET-16 topology, the amount of protection bandwidth for path protection is a little more than the working bandwidth. This is comparable to the amount of protection bandwidth in SONET shared protection rings.

REFERENCES

- [1] N.G. Duffield, P. Goyal, A. Greenberg, P. Mishra, K.K. Ramakrishnan, and J.E. van der Merwe, "A flexible model for resource management in virtual private networks," *Proceedings ACM SIGCOMM 1998*.
- [2] A. Kumar, R. Rastogi, A. Silberschatz, and B. Yener, "Algorithms for provisioning virtual private networks in the hose model," *Proceedings ACM SIGCOMM 2001*.
- [3] G. Italiano, R. Rastogi, and B. Yener, "Restoration algorithms for virtual private networks in the hose model," *IEEE INFOCOM 2002*.

¹This work was supported by Fujitsu Laboratories of America.