

# Trade-offs in Resource Management for Virtual Private Networks

Satish Raghunath  
Nortel Networks, MA, USA  
Email: rsatish@alum.rpi.edu

Shivkumar Kalyanaraman  
Department of ECSE, RPI, NY, USA  
Email: shivkuma@ecse.rpi.edu

K.K. Ramakrishnan  
AT&T Labs - Research, NJ, USA  
Email: kkrama@research.att.com

**Abstract**—Virtual Private Networks (VPNs) feature notable characteristics in structure and traffic patterns that can be exploited by the service provider to achieve significant capacity savings.

Efficient provisioning of point-to-point connections using statistical admission control is well understood. However, provisioning a VPN involves provisioning a set of point-to-multipoint connections and features an additional dimension in the form of a traffic matrix. Consequently we have multiple network mechanisms that are important for efficient operation - a) admission control, b) signaling-based per-link reservations, c) traffic matrix estimation. In this paper we examine the relative importance of mechanisms that positively affect the operational efficiency in the context of VPN provisioning.

Using insights from our extensive measurement based study on the structural properties usually observed in VPNs, we build a simulation framework to quantify the trade-offs in opting for one mechanism over the other. We arrive at our conclusions with the help of simulations featuring a variety of VPN structures and network topologies. We find that the structural characteristics of VPNs cause traffic matrix estimation to be a dominant factor in determining the utilization gains. Consequently, we find that deploying statistical techniques might not be worth the effort if the traffic matrix is not incorporated. While signaling-based reservation mechanisms lead to higher utilization, edge-based techniques prove to be lot more scalable and simpler to realize. We explore the means to reduce the performance penalty associated with such simpler techniques.

## I. INTRODUCTION

A Virtual Private Network (VPN) securely connects multiple customer sites that wish to communicate among each other. The motivation for customers to use a VPN is often the service assurances obtained from the provider in terms of a pre-specified Quality of Service assurance. The Service Level Agreement (SLA) is in the form of assured bandwidth, expected loss rates and delays. A service provider then provisions the network to ensure that the SLAs for an admitted VPN are met based on information provided by the VPN customer.

The QoS achievable for a given VPN is influenced by the way customer sites are inter-connected by the provider. The most straightforward solution is to have a mesh of point-to-point links connecting customer sites (Fig. 1). A more efficient and scalable solution would be to *multiplex* multiple VPN customers on a common core network (Fig. 2). In such a network, the provider can obtain high resource utilization and simultaneously ensure SLAs with the help of adaptive network mechanisms that exploit statistical properties of customer traffic, including admission control, queuing and scheduling.




Fig. 1. Virtual Private Networks can be provisioned using a mesh of point-to-point links




Fig. 2. A scalable strategy to provision VPNs uses common multiplexed core network

When connectivity requirements are point-to-point in nature, the provider may achieve efficient network operation using a rich set of statistical techniques to characterize traffic and perform admission control [5]. These techniques have been evaluated in a wide range of scenarios and their gains over deterministic admission techniques have been quantified [9].

However, when one considers provisioning a network for VPN customers, as we shall observe later in this paper, a strategy that recognizes and takes advantage of the distinct features of VPNs may outperform one that does not, even if per-hop statistical admission control techniques are employed. This is due to the fact that provisioning a new VPN involves a set of *point-to-multipoint* requests. Admitting and provisioning the VPN requires the entire set of endpoints of the VPN to be admitted together. Understanding the traffic matrix for the VPN is useful because, in addition to describing the traffic demands, a traffic matrix also depicts the structure of interactions in the VPN.

Consider for example, the VPN with endpoints  $E_1, E_2, E_3, E_4$ , as shown in Fig. 2. The provider edge (PE) routers corresponding to these endpoints are denoted as

$P_1, P_2, P_3, P_4$ . The aggregate traffic arriving into the network at  $E_1$  (referred to as  $T_1$ ) possibly contains traffic toward  $E_2, E_3$  and  $E_4$ . Provisioning bandwidth for the VPN involves provisioning for the aggregate from endpoint  $E_1$  as well as from the other endpoints  $E_2, E_3$  and  $E_4$ . If we consider  $T_1$ , one option would be to assume that the aggregate demand from  $E_1$  could be directed toward any of the other endpoints. Alternately, a traffic matrix estimating the portion of demand that is directed toward each of  $E_2, E_3, E_4$  could be employed in addition to the aggregate specification of  $T_1$ . The admission decision then involves checking if the capacity available can support all the aggregates from all the endpoints of the VPN.

While intuition suggests that incorporating traffic matrix information in provisioning may deliver higher efficiencies, the question is if there are enough gains in the context of VPNs, to justify deploying a measurement infrastructure to obtain such information. In addition, we would want to know if comparable performance gains can be obtained by deploying an improved admission control algorithm or signaling-based reservations. In essence, the question is one of quantifying the relative importance of multiple means to achieve higher performance in provisioning VPNs. This is an important question for providers because, in reality, obtaining and analyzing per-customer and per-hop information involves much effort. The per-destination traffic characteristics (the traffic matrices) are usually not known *a priori* and require algorithms that learn such information [13]. On the other hand, if traffic matrices are measured after admitting VPNs using an initial coarse estimate, the admission control process can be continually refined to better reflect available resources. Such refinement may involve exploiting statistical admission control techniques or simply better deterministic estimates of demand. Similarly, signaling-based mechanisms imply complexity in administration and management.

We find that studying this question with reference to properties of VPNs leads to new and important insights. Using recent findings from our extensive measurement study [13] of SNMP data from a large IP/VPN service provider, we capture VPN structural properties in a simulation framework. Typically, VPNs feature distinct structural characteristics. E.g., many VPNs fall in the category of a hub-and-spoke VPN. In such a VPN, there is a customer site which acts as the hub for all communication in the network and the communication from the other sites in the network is primarily to and from this hub site. Clearly, while provisioning such a VPN the admission control mechanism needs to primarily consider the path to the hub site. A meshed VPN with all of the sites being peers of each other results in multiple paths being utilized with a more complex admission control decision needed.

We present results from a large number of experiments using several commercial backbone topologies. The important findings of our study include: a) A technique exploiting VPN structure can yield significant gains over another that does not, even when we employ signaling based statistical admission; b) Even modest increase in the complexity of VPN interactions dramatically increases the gains of using traffic matrices;

c) Scalable alternatives to signaling-based architectures can considerably reduce performance penalties by incorporating a combination of traffic matrix information and dynamic resizing of allocations.

The rest of paper is organized as follows. In §II we discuss related work and the context in which our contributions apply. We define the various parameters and strategies along with the framework in which we evaluate them in §III. In §IV we briefly present the results of a measurement study that were the basis for the simulations. Results are presented in §V and §VI. We summarize in §VII.

## II. RELATED WORK

This paper relates to two broad areas of research, namely admission control and VPN resource allocation models.

The goal of admission control is to provide a pre-specified level of QoS to flows in the network. There have been several admission control proposals aimed at providing statistical and deterministic QoS assurances (see e.g., [5], [9]). Statistical schemes are able to exploit the bursty nature of multiplexed traffic to deliver much higher gains. Typically, these proposals derive a statistical characterization of the sources and compute the probability of violating the QoS parameters. Recent work has concentrated on building statistical assurances using simple deterministic components [1], [8], [14] like leaky-bucket shapers. We employ such a statistical admission scheme in succeeding sections.

The question of how to provision Virtual Private Networks has been of interest to many researchers recently. If the demands are assumed to be known before-hand, the problem can be solved for those demands using optimization techniques [6], [7], [10]. In contrast, the Hose Model [3] provides a solution that exploits multiplexing gains by admitting VPNs on a common core network without needing traffic demand matrix information, while relying on signaling based reservation and adaptive resizing. A more recent solution to the problem is provided by the Point-to-Set model [11], [12] where the authors employ traffic demand matrix estimation techniques while not requiring signaling and not assuming that demands are known a-priori.

These proposals are complementary to the present treatment. We quantify the importance of network mechanisms like signaling and the gains that traffic matrix information can yield. We examine these issues in the context of VPN structure and admission control strategy. Thus our work places these solutions in the larger context of VPN resource allocation trade-offs. If a designer can weigh the importance of each mechanism to a specific implementation situation, the appropriate model can be easily picked.

## III. MECHANISMS AND PARAMETERS OF INTEREST

We begin our study with a discussion of the parameters that affect scalability and achievable resource utilization gains. A range of resource management solutions can be conceived by exploiting various combinations of the following properties.




Fig. 3. (a) A Hub/Spoke VPN; (b) A VPN with each endpoint communicating with multiple endpoints,  $\eta$  is the maximum number of endpoints with which a given endpoint communicates

### A. VPN Structure

Based on an extensive measurement study of a large IP VPN service using several months of data, we identified the spatial and temporal characteristics of customer VPNs [13]. We identify three broad categories of VPNs: pure hub/spoke, meshed and hybrid VPNs. We observe that a commonly occurring structure in VPNs is the “Hub/Spoke” model (approximately 48% in the measurement study, briefly described in §IV). As seen in Fig. 3(a), there is a central “Hub” node with which every other endpoint communicates. As far as the provider network is concerned, the capacity requirements are between a given spoke and the hub.

We cover other structures using a generic model (seen in Fig. 3(b)). This model features a parameter  $\eta$  indicating the maximum number of endpoints with which an endpoint communicates in the VPN. In Fig. 3(b) one of the endpoints has four arrows leading from it indicating its communicating peers. While studying the impact of the structure of VPNs on admission control strategies, we vary: a) the percentage of VPNs that belong to each category and b) the value of  $\eta$  for generic VPNs.

### B. Admission Control

Given parameters characterizing the traffic, the provider has to make a decision of whether to admit a new customer. In the succeeding sections we use the dual leaky-bucket specification to describe the traffic from a given customer endpoint. The dual leaky-bucket consists of three parameters  $(\pi, \rho, \sigma)$  indicating respectively the peak rate, sustainable average rate (leaky-bucket rate) and the burst parameter (bucket size). If the arrival process is indicated by  $A(t)$ , conformance to  $(\pi, \rho, \sigma)$  means  $A(t, t + \tau) \leq \min(\pi\tau, \rho\tau + \sigma)$ . Thus for an endpoint  $i$ , the traffic specification is given by  $(\pi_i, \rho_i, \sigma_i)$ . With this data, one could opt for either a probabilistic admission test or a deterministic one.

1) *Statistical Admission Control*: Statistical admission schemes typically evaluate the probability of violating a given QoS metric. If loss is the chosen metric, the admission test would be to ensure that the new flow being let into the system does not increase the probability of loss beyond a particular threshold. If  $L$  denotes the random variable for loss, the

condition could be  $Pr\{L > 0\} \leq \epsilon$  where  $\epsilon \in (0, 1)$  is a pre-specified parameter.

There have been numerous proposals for statistical admission control tailored for various assumptions regarding traffic information and network topology. Statistical multiplexing of sources on a link helps in achieving high resource utilization. However, multiplexing alters statistical properties of the flow. In our case, the need to couple statistical admission with signaling means that such computations must take into account distortion effects introduced by multiplexing. In general, it is hard to quantify the degradation in QoS due to multiplexing across multiple hops [5]. The options would then be: a) introduce per-hop shaping to eliminate distortions due to multiplexing and use one of the several proposed schemes designed for a single link; b) Opt for a scheme which avoids per-hop shaping and still ensures mathematical tractability for per-hop computations.

The first option involves configuring shapers at each hop to retain properties of the flow. To avoid losing focus to per-hop shaper parameter setting issues, we chose the second option. The recent statistical QoS framework by Reisslein et al [14] provides such a solution. They adopt bufferless multiplexing so that the independence assumption among flows is valid inside the network and hence computations are far simpler. They demonstrate that, with the right smoothing operation at the entry of the network, their scheme performs equally well as buffered statistical schemes (e.g. [4]). Please refer to Appendix I for details.

While these techniques have been applied toward admission control for point-to-point requests, we adapt them for the context of VPNs where the requests are actually a set of *point-to-multipoint* connections. In order to do that, we build a model of the point-to-multipoint connection using traffic matrix information so that the equivalent set of point-to-point requests can be deduced. We shall examine this in further detail in §III-D and Appendix I.

2) *Deterministic Admission Control*: In our context, where each endpoint of the customer VPN generates a point-to-multipoint flow, the primary customer specification is the rate of this flow coming into the network. The customer VPN is a set of such point-to-multipoint flows which have to be admitted together when the VPN customer is accepted. A simple strategy in admitting a new flow is to quantify its peak bandwidth requirements and reserve that capacity inside the network. Thus, a first option would be to reserve the peak rate specified by the customer. However, if more elaborate traffic matrix information is available apriori, e.g., mean and variance of expected load on a source-destination pair basis, we could enhance this scheme. The reservation could then be equal to the mean in addition to a multiple of the standard deviation.

### C. Signaling

In the presence of network support, admission control and bandwidth reservation decisions can incorporate information from each hop of a path along which a flow is admitted, as

is typically seen in signaling based admission control. Given the traffic characteristics at each hop, we could either perform statistical or deterministic admission control and exploit traffic matrix information, if it is available apriori. Such a framework allows for high resource utilization to be achieved, but relies on a lot information and support from both the customers and the network. While the resource gains are desirable, we would certainly wish to relax the amount of network support required. We observe that the evolution of deployed mechanisms in the network has avoided complex signaling and admission control protocols for QoS. That leads us to the question of what can be done in the absence of support for signaling. If admission decisions should not involve per-hop computations we can think of the following options:

1) *Centralized Admission*: A Centralized admission control entity can have up-to-date knowledge of the whole network and hence achieve the same efficiency as a signaling-based approach. The disadvantages of this, however, would be that there is now a single point-of-failure, and the centralized entity has to maintain state about the whole network. It is preferable to have the admission test carried out in a distributed fashion.

2) *Distributed scheme with fixed path capacities*: Signaling-based reservation prevents the race condition where a shared resource might be over-booked by multiple users. In addition to dealing with the race condition we typically observe in the case of admission control of multiple independent *point-to-point* flows, our distributed scheme has to deal with two additional complications: that of dealing with the *point-to-multipoint* flow from each VPN endpoint; that of ensuring that all such flows from all the endpoints of a customer VPN are considered and admitted *simultaneously*. Further, a distributed scheme without signaling would have to feature admission control decisions at each entry-point in the network independently of other network edges.

To solve this problem, we could evolve some means of apportioning capacity among the source-destination pairs so that they can make decisions independently. If the topology and routes between source-destination pairs are known apriori, we could build algorithms which assign capacities to paths. Algorithm 1 represents one such scheme.

The ability to independently make admission decisions at each network entry point is clearly an advantage we wish to have. The disadvantage with the approach is that it is oblivious to traffic trends. If some paths carry less traffic, that share of the capacity would be wasted. We solve these issues with an improved approach in §VI using an architecture for dynamic adaptation of assigned path capacities.

#### D. Traffic Matrix Information

The final parameter we introduce is the customer traffic matrix. The traffic matrix specifies information about expected traffic between a given pair of endpoints in the customer VPN.

Typically, the traffic originating from a source node is split among a set of egresses. If there is information about per-destination traffic trends, the allocation can be tailored accordingly. If the source rarely directs traffic at the peak

---

#### Algorithm 1 Path Capacity of path $p$

---

Denote capacity of link  $l$  as  $C_l$  and that of path  $p$  as  $C_p$   
Input:  $\mathcal{L}_p \leftarrow \{ \text{Set of links in } p \}$   
Input:  $\mathcal{P}_l \leftarrow \{ \text{Set of paths traversing link } l \}$   
**for** each link  $l \in \mathcal{L}_p$  **do**  
     $|P_l| \leftarrow \text{Number of paths traversing link } l$   
     $S_p(l) \leftarrow \frac{C_l}{|P_l|}$   
**end for**  
 $C_p \leftarrow \min_{l \in \mathcal{L}_p} S_p(l)$

---

rate (e.g., based on the access bandwidth) toward a single destination, there can be multiplexing gains compared to peak provisioning. Information about pair-wise traffic trends could either be specified by the customer or some measurement-based mechanism may be used to learn these trends. In the absence of such pair-wise information about the traffic matrix, there would be over-provisioning of the network links. But the upside would be a simpler framework.

As introduced in §III-B, we employ a  $(\pi, \rho, \sigma)$  specification to describe the aggregate traffic. The traffic matrix could then be specified either as:

- 1) A set of triples  $(\pi_j, \rho_j, \sigma_j)$  governing the traffic toward endpoint  $j$  so that  $\sum_j \rho_j = \rho$ ,  $\sum_j \sigma_j = \sigma$  and  $\pi_j \leq \pi$  or,
- 2) A set of mean and variance values  $(m_j, v_j)$  for random variables  $p_j \in [0, 1]$  which represent the fraction of the aggregate directed toward destination  $j$ . Thus if  $A$  is the aggregate traffic and  $A_j$  is toward destination  $j$ , we have  $A_j = p_j A$  and  $A = \sum_j A_j$

We choose the second option since it is a more intuitive description (e.g., 70% of the aggregate is directed toward destination 1), and leads to convenient implementation. It can be shown that a dual leaky-bucket description for per-destination traffic can be deduced from  $(m_j, v_j)$  and  $(\pi, \rho, \sigma)$  (please see Appendix II for details).

#### IV. INSIGHTS FROM MEASUREMENT INFORMATION

The applicability of the simulation study depends on the feasibility of the assumed models to generate VPNs, perform admission control and represent traffic information.

The simulation framework employed in the following sections exploits the inferences of an extensive measurement study of a large IP/VPN service provider that the authors conducted recently [13]. Similarly, the choices discussed in the previous sections are motivated by this measurement study. We briefly present the aspects related to the current simulations that are derived from that study.

Examining the properties of IP/VPNs using entropy-based traffic matrix estimation techniques applied to SNMP measurement data yields important hints about underlying nature of VPNs. One of the most important observation is that a large number of VPNs fall in the category of Hub/Spoke VPNs where there is central communication hub node with which all other endpoints communicate. A classification of VPNs leads




Fig. 4. Results of an extensive measurement study [13] showed hub/spoke structures are the most common among IP VPNs




Fig. 5. CDF of number of Endpoints per VPN

to the pie-chart depicted in Fig. 4. Consequently, we accord special importance to this structure in the simulations.

We also found that the temporal characteristics of VPN traffic feature interesting properties. They demonstrate stable trends across multiple weeks and in some instances even months. The traffic observed from a particular VPN endpoint toward other endpoints of the VPN is observed to again demonstrate such stable trends and is amenable to estimation and learning techniques. This means modeling VPN traffic matrices with per-destination random variables with a mean and variance is feasible. This is the approach we use in quantifying traffic matrix information.

In order to generate VPNs synthetically, we employ empirical distribution of sizes of VPNs (obtained from the previous measurement study) in the simulation results presented in future sections. We reproduce the cumulative distribution function of VPN sizes in Fig. 5.

## V. COMPARATIVE ANALYSIS

In the following paragraphs we proceed by considering different interesting combinations of the mechanisms described in the previous sections.

### A. Topologies and experimental setup

We use topologies of popular commercial backbones (namely, AT&T, Sprint, Qwest and MCI, see Table I) for our experiments [2]. Each experiment consists of two phases - a VPN generation phase and an admission control phase. The

Network	Links	Nodes
AT&T Worldnet	144	86
Sprint	71	47
Qwest	50	17
MCI	31	19

TABLE I  
TOPOLOGIES EMPLOYED IN SIMULATIONS

experiment is started with a set of values for the structure of VPNs to be generated and dual-leaky-bucket parameters for the aggregate traffic from an endpoint. The VPN generation routine then produces VPNs of varying sizes using the empirical distribution from measurements (Fig. 5). These VPNs are then fed to the admission control routine one after another. Since the VPNs are generated using a randomized procedure (described below) for *each* experiment and due to the fact that a large number of experiments are conducted, the results do not depend on the order in which the VPNs are admitted.

To generate a VPN endpoint's characteristics randomly the following procedure was followed. In the case of the a Hub/Spoke VPN, there is no need to generate the set of peers for each endpoint. Otherwise, every endpoint in the VPN has a set of destinations with which it communicates. The procedure involves picking a random subset of endpoints and deciding the fraction of traffic that is destined to each destination. E.g., consider an endpoint with a destination set with  $K$  nodes;  $(K - 1)$  uniform random numbers,  $r_i, i = 2 \dots K$  are generated in the range  $[min, max]$ . Recall that  $m_j$  is the mean of the per destination fraction of the traffic to destination  $j$  from the endpoint under consideration. Then setting  $r_1 = 1$  and  $m_1 \sum_{i=1}^K r_i = 1$ , we obtain  $m_i = r_i m_1$ . The variance of per-destination traffic fraction  $v_j$  is then computed as a fixed fraction of the mean. Choosing the variance to be higher only reduces the improvement delivered by statistical admission techniques, but does not affect the utility of traffic matrix information. The reader is directed to the measurement study [13] for more on temporal traffic trends in VPNs. The range  $[min, max]$  decides the *bias* toward a subset of destinations in the set. If the range is small and around 1, traffic is equably directed to all nodes in the set. Higher the value of  $max$  greater the spread of the load distribution among destinations. The dual-leaky-bucket regulator parameters for all VPNs was set at  $(0.5 Mbps, 0.15 Mbps, 20 kb)$ . The link capacities are available in the topology data files [2] available online except in the case of the MCI topology, where the link capacities were set to  $100 Mbps$  and their delay was chosen to be  $10 ms$ . Except when specifically mentioned, the results shown are limited to those from the MCI topology due to space limitations.

In the following sections we study the role of the mechanisms introduced in §III. We primarily examine the change in the number of customer VPNs accepted when a mechanism is used in the decision process.




Fig. 6. Number of Admitted VPNs in the presence of signaling-based per-hop admission control with 30% of the generated VPNs being of the Hub/Spoke type

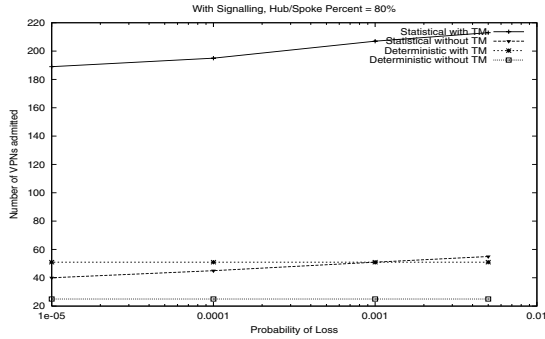


Fig. 7. Number of Admitted VPNs in the presence of signaling-based per-hop admission control with 80% of the generated VPNs being of the Hub/Spoke type

## B. Traffic Matrix

In §III-D we specified traffic matrix information in terms of the mean and variance of the random variable representing the per-destination traffic fraction. Thus, if  $A$  were the aggregate traffic from an endpoint and  $A_j$  were the traffic directed toward destination  $j$ , we introduced a random variable  $p_j \in [0, 1]$  so that  $A_j = p_j A$ . make

In the presence of traffic matrix information, the admission control decision for a link need only account for the fraction of traffic that is likely to be directed along this link. In particular, we evaluate the admission criterion considering a dual leaky-bucket specification derived using  $(m_j, v_j)$  and  $(\pi, \rho, \sigma)$  (please see Appendix II). In the absence of such information, the admission decision assumes  $(\pi, \rho, \sigma)$  as the specification of traffic toward every destination. In discussions in the rest of this subsection, simulations feature signaling-based admission control and examine the utility of traffic matrix information. The plots in Figures (6) and (7) show the benefits of collecting traffic matrix information. The salient points to be noted are:

- The cost of not exploiting VPN structure is significant. In terms of admission control, the best one can do is a statistical per-hop process. The figures indicate that over and above such a mechanism, VPN structure information delivers dramatic gains. E.g. in Figure (6) the plot indicating results for statistical admission control featuring




Fig. 8. The utility of statistical admission control reduces when compared to deterministic admission in the absence of traffic matrix

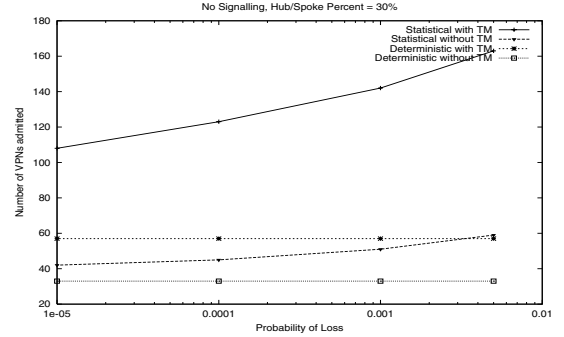


Fig. 9. Number of admitted VPNs falls in the absence of signaling-based admission control (percentage Hub/Spoke VPNs = 30%)

traffic matrix information shows a significant gain over the plot obtained without traffic matrix information.

- A statistical admission control algorithm without traffic matrix information does as well as a conservative admission control scheme that exploits traffic matrix information (compare the plots marked “Statistical without TM” and “Deterministic with TM”). Intuitively, with more information about traffic matrix and VPN structure, the admission scheme can be simpler for the same resource utilization gain.

Thus if a majority of the VPNs being serviced are of the Hub/Spoke nature and no traffic matrix information is available, a simple deterministic admission can be a good choice if it can be enhanced with information about what nodes are spokes and which node is a hub. Fig. 8 illustrates this reduction in gain of a statistical scheme over a deterministic with increasing fraction of VPNs being of the Hub/Spoke kind.

## C. Signaling-based admission

In §V-B we looked at the benefit of traffic matrix information when used in conjunction with signaling-based admission control. We now consider the situation where traffic matrix information is available and compare the gains that may be achieved with and without signaling-based admission control.

In the absence of signaling-based admission control, we have to make admission control decisions at the entry of the network. In §III-C we discussed the options available in the absence of signaling. We proposed a simple and static

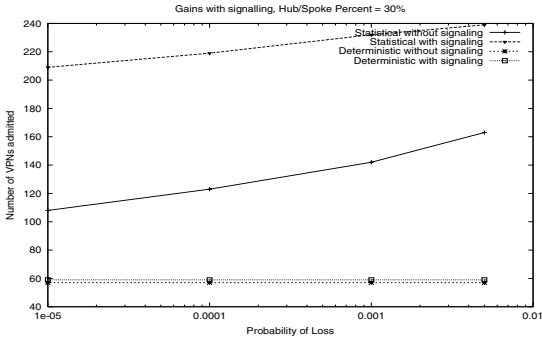


Fig. 10. Gains in the number of admitted VPNs with Signaling in addition to traffic matrix information

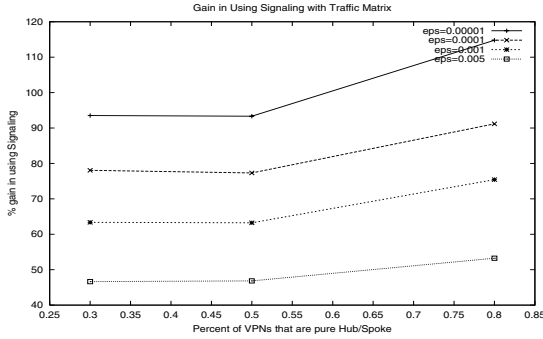


Fig. 11. Signaling-based admission control is superior irrespective of the percentage of Hub/Spoke VPNs (with traffic matrix)

path capacity computation algorithm so that admission control decisions are made considering the ingress-to-egress path as a virtual link with capacity derived from Algorithm 1. We now employ this strategy to evaluate the value added by using signaling-based admission mechanisms.

Fig. 9 shows the result of such an edge-based strategy for different loss probabilities. Comparing this with Fig. 6 we can clearly see the reduction in the number of admitted VPNs. The plot in Fig. 10 confirms this inference and shows the gains with signaling. Figures (11) and (12) present this aspect across varying nature of generated VPNs. The following observations are in order:

- The trends indicate that signaling yields consistent gains irrespective of the structure of the VPNs and the availability of traffic matrix.
- While the path capacity algorithm is simple and enables edge-based admission decisions, it does not perform as well as an algorithm that exploits signaling.

We would certainly want to retain the simplicity of the edge-based admission scheme while obtaining the performance comparable to a signaling-based mechanism. In §VI we examine strategies to bridge the gap in performance via an improved algorithm.

#### D. Effect of structure of VPNs

We examine the effect of structural properties of VPNs using two metrics: a) the percentage of VPNs that are

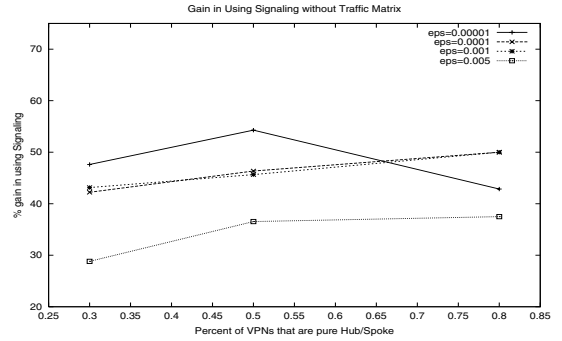


Fig. 12. Even in the absence of traffic matrix information Signaling-based admission control is superior irrespective of the percentage of Hub/Spoke VPNs

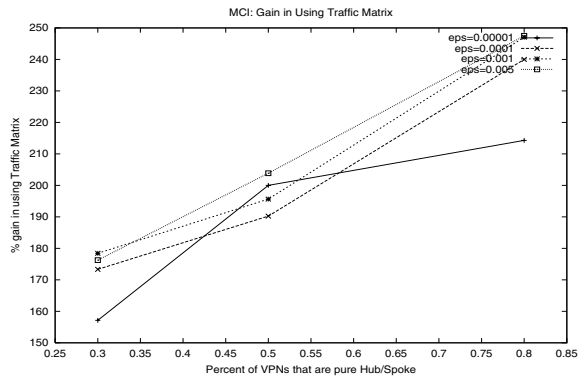
Hub/Spoke; b) a more generic parameter  $\eta$  denoting the number of communicating peers for each VPN endpoint.

Since the structure of the VPN can be best captured by measuring and exploiting its traffic matrix, we shall examine the gains in terms of the number of admitted VPNs if traffic matrix information was available. Thus Fig. 13 shows the percentage gains with statistical admission if traffic matrix information is introduced. Although one would expect to improve admission gains with traffic matrix information, what is noteworthy is that the quantum of gains can be very significant. The variations in the quantum of gains with the Hub/Spoke percentage does not seem to be independent of topology as indicated by Fig. 13.

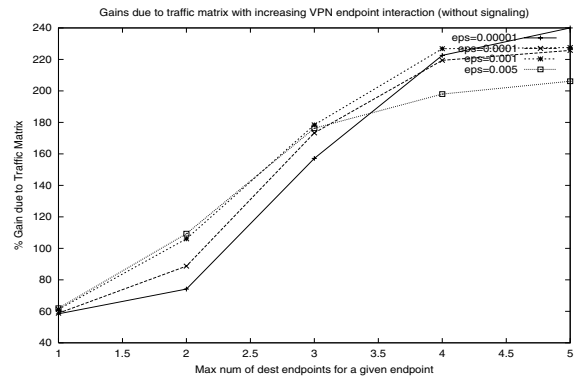
We now examine the role of  $\eta$  in more detail. A VPN with higher  $\eta$  has more complex interactions among its endpoints. Our experiments support the intuition that understanding the traffic matrix when  $\eta$  is higher would yield higher resource utilization gains.

Fig. 14 depicts the importance of traffic matrix with increasing complexity in VPN endpoint interactions. The service provider can get a significant benefit in terms of resource utilization (due to the ability to admit a larger number of VPNs) by taking advantage of the traffic matrix characteristics. This is particularly true as we go away from a simple hub and spoke VPN structure (when  $\eta = 1$ ) to a VPN with more peer-to-peer communication. Similarly signaling gains (Fig. 15(a) and Fig. 15(b)) become significant with higher  $\eta$  when there is no traffic matrix information. The results presented till now confirmed and quantified the intuition that complicated VPN structures imply significant costs in resource allocation if we do not take advantage of the benefits of the traffic matrix or signaling. Fortunately, we can devise strategies to exploit such information without elaborate changes to the network:

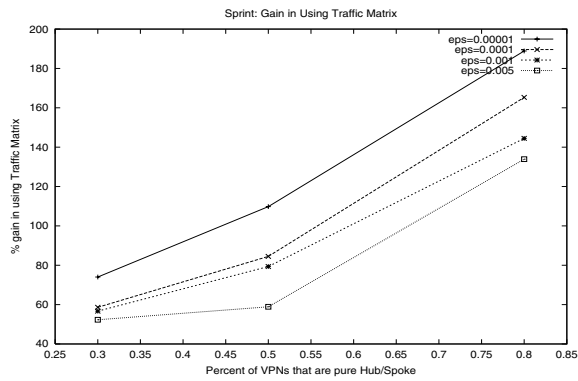
- 1) Recent research (e.g., [13], [15]) has led to efficient means of estimating large traffic matrices using long-term SNMP link statistics. This implies that the admission control process can be adapted and refined continually as we learn more about the traffic matrices of all previously admitted VPNs.
- 2) Improved algorithms to manage edge-based path capacity allocations dynamically (as discussed in §III-C) can lead to performance that compares well with signaling-



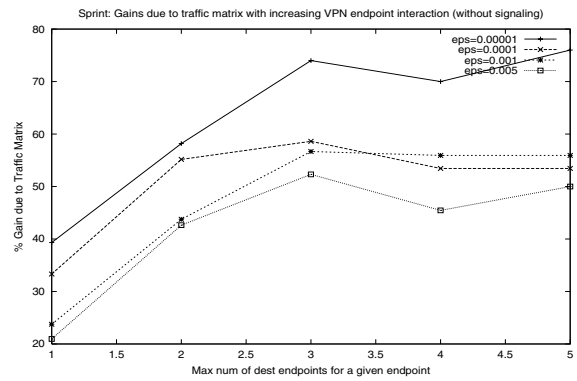
(a) MCI



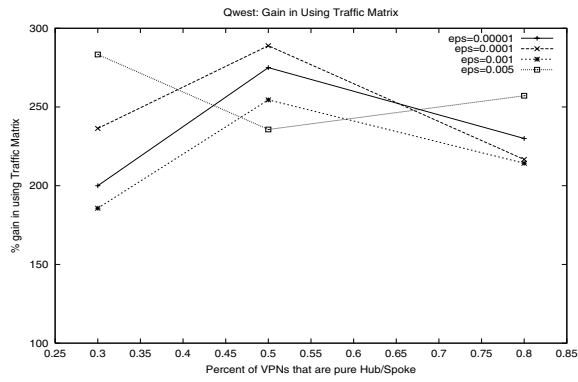
(a) MCI



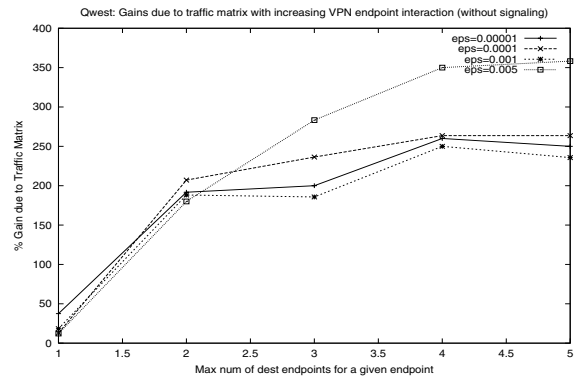
(b) Sprint



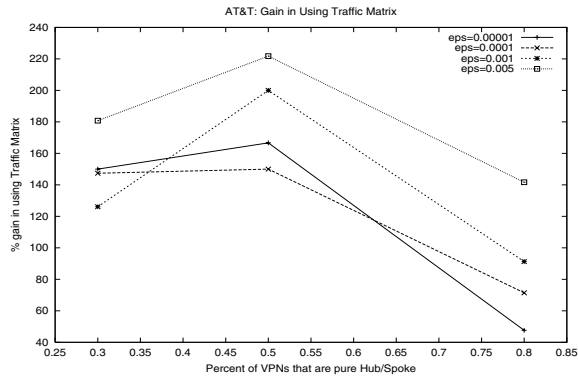
(b) Sprint



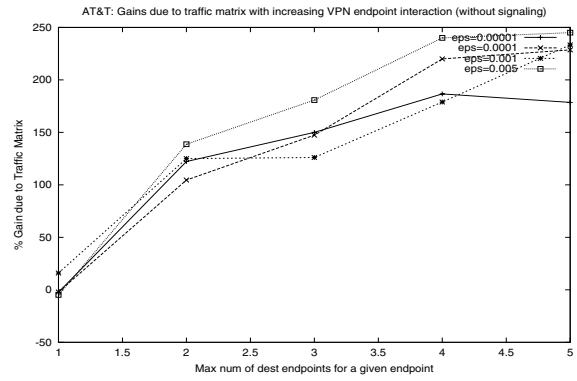
(c) Qwest



(c) Qwest



(d) AT&T



(d) AT&T

Fig. 13. Using Statistical Admission Control: although incorporating Traffic Matrix consistently provides admission gains with varying number of Hub/Spoke VPNs, the quantum of gains depends on the specific topology.

Fig. 14. With increase in the number of endpoints with which a node communicates (higher value of  $\eta$ ) gains due to traffic matrix become more pronounced.

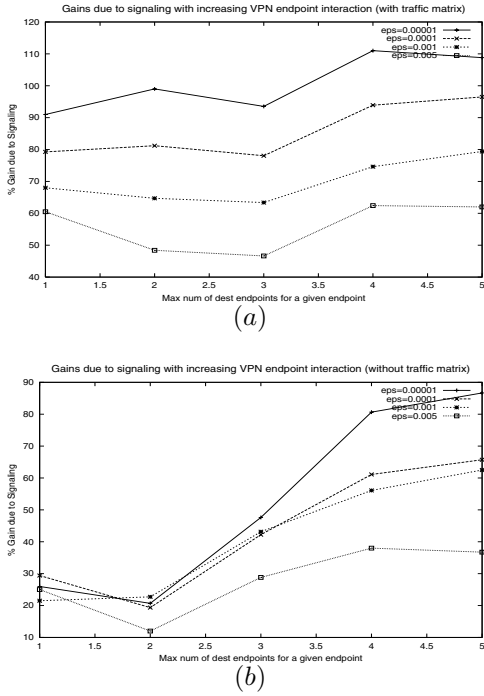


Fig. 15. (a) With traffic matrix available, gains due to signaling hold steady across multiple values of  $\eta$ ; (b) In the absence of traffic matrix, signaling gains become significant as  $\eta$  grows

based admission control schemes.

We elaborate on the second point in the following section by improving the path capacity algorithm presented in Algorithm 1. In summary, we examined the different parameters that affect resource utilization and quantified their importance. We found that learning more about the nature of VPNs to be served (e.g., are they hub and spoke) allows us to exploit attractive trade-offs (e.g., simple deterministic schemes in presence of majority hub and spoke type VPNs).

## VI. DYNAMIC PATH CAPACITY

In §III-C we introduced a simple path capacity computation algorithm in order to substitute the function of signaling-based admission. The algorithm statically divided link capacities among different source-destination paths traversing the link. There are some notable disadvantages to this algorithm: a) It does not consider the fact that some paths carry more traffic than others. A static link sharing scheme does not reassign bandwidth to other paths which might be seeing higher demand; b) It assumes that routing and topology are fixed. The capacities are computed assuming that the links along a source-destination path are known; c) It requires the network edge to process routing control information and compute path details. In this section we attempt to remedy these drawbacks and describe an improved algorithm. In doing so, we increase the resource allocation gains while avoiding signaling-based mechanisms.




Fig. 16. The network edges can be decoupled from routing and topology changes if they communicate a central measurement server which provides path capacity information

### A. Distributed Admission, Centralized Measurement

In order to remedy the drawbacks of the aforementioned algorithm, we envisage decoupling the functions of computing the path capacities and making the admission control decision. The former involves processing routing information for topology and capacity details. The latter involves computing an admission control criterion given traffic characteristics. Thus these are separable tasks. Fig. 16 demonstrates such an architecture. A central “measurement server” receives routing updates so that it has a snapshot of the topology. It processes this data to compute path capacity values for all ingress-egress pairs. The network edge evaluates the admission test using the capacity information obtained from this central server.

The clear advantage with this setup is that routing and topology changes are shielded from the network edges which make admission decisions. The edge is periodically notified of the path capacity that is available toward any destination edge it might want to reach.

Further, the admission process can be continually refined with better estimates of traffic demand characteristics for existing VPNs. If the provider chooses to estimate traffic matrices, this can be achieved by having the admission control algorithm utilize per-destination information when it becomes available.

With this architecture in mind we devise an improved path capacity assignment algorithm (Algorithm 2). This algorithm uses a parameter  $\beta \in [0, 1]$  which indicates the fraction of link capacities that are statically pre-assigned to each path according to Algorithm 1. Lower the value of  $\beta$  higher the flexibility to the allocation algorithm in assigning path capacities. If the value of  $\beta$  is too low, it might cause some network edges to refuse admission even when there is capacity available. Thus this parameter needs to be tuned to obtain acceptable behavior.

The algorithm at the network edge now becomes much simpler. As specified in Algorithm 3 it evaluates the admission criterion to arrive at a probability of loss. If the probability is less than a pre-determined threshold  $\epsilon$ , the request is admitted. Further, if the probability of loss is within a factor,  $\alpha$ , of the threshold (i.e., we are low on available capacity) or if the admission control test fails, it requests for additional capacity

---

**Algorithm 2** Dynamic Path Capacity Computation

---

Precondition: Apportion  $\beta C_l, \beta \in [0, 1]$  equally among all ingress-egress pairs.

Input: Current routing and topology state

Input: Request for additional path capacity

Input:  $C^*$  representing a capacity increment.

**if** Unused capacity exists **then**

    Accept request and increase source-destination path capacity by  $C^*$

**end if**

---

from the central measurement server.

---

**Algorithm 3** Admission Control at an edge

---

Input: A point-to-multipoint service request from an endpoint  $E_1$  toward a set of egresses  $P_1, P_2, \dots, P_n$

Input: Capacity available on path between  $E_1$  and  $P_i$  obtained from Measurement Server

Input:  $\alpha \in (0, 1)$  decides when to request for more capacity.

**for** each path  $(E_1, P_i)$  **do**

    Compute probability of loss  $P_{loss}$

**if**  $P_{loss} > \epsilon$  **then**

        Request measurement server for additional path capacity

        Reject admission request

        Return

**end if**

**if**  $P_{loss} > \alpha\epsilon$  **then**

        Request measurement server for additional path capacity

**end if**

**end for**

Accept admission request

---

### B. Results

We now evaluate the algorithms presented in the previous section. Fig. 17 demonstrates the gains in dynamically apportioning bandwidth versus a static algorithm. For these experiments we set  $\alpha = 0.01$ . The value of  $\alpha$  decides how soon a path capacity is reassigned with reference to the time the SLA may be violated. In our experiments varying  $\alpha$  over a range within the same order of magnitude did not affect the results significantly. Next we present, in Fig. 18, the effect of varying  $\beta$  from 0.2 to 1.0 (equivalent to the static apportioning algorithm). As expected reducing  $\beta$  provides more flexibility in allocating path capacity and allows for higher number of admitted VPNs.

In summary, the improved path capacity algorithm provides for higher gain and compares better with signaling-based mechanisms as compared to the static path sharing scheme.

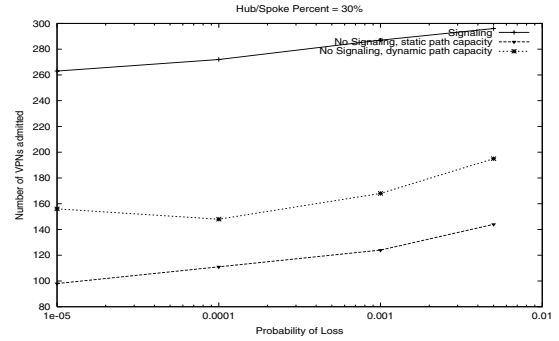


Fig. 17. The Dynamic path capacity allocation considerably improves the performance of the static link sharing scheme

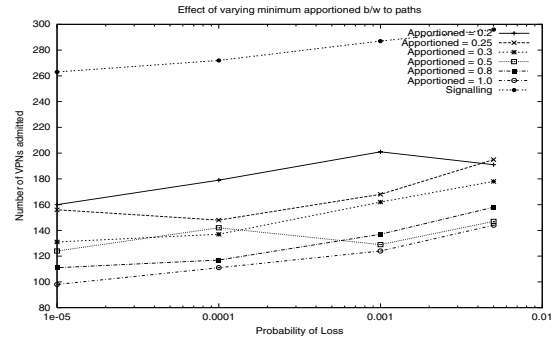


Fig. 18. With lower values of  $\beta$  we have more flexibility in allocating path capacity where there is demand and hence more gain

## VII. SUMMARY AND CONCLUSIONS

In this paper we examined all the mechanisms that influence resource allocation in Virtual Private Networks and quantified their relative importance. We presented a set of mechanisms and parameters available to a service provider that affect VPN provisioning and the achievable resource utilization; viz., the admission control strategy, the availability of traffic matrix, the information about VPN structure and support for signaling based admission control. We also looked at the various options available to the service provider with each of these mechanisms.

We parameterized our simulation framework based on an extensive measurement study on the structural properties usually observed in VPNs. We then conducted experiments with four different backbone topologies, and considered statistical and deterministic admission control strategies in a variety of scenarios. Our experiments to understand the interplay of these factors led to several important conclusions:

- When the traffic matrix information is available, it has a dominant effect on the resource utilization gains. Knowledge of the structure of the VPN becomes important since it has an influence on the traffic matrix (e.g., knowing that a VPN is of the Hub/Spoke type implies knowing most of the traffic matrix).
- Signaling-based admission control can vastly improve resource utilization. However, in the absence of signaling, the penalty of simpler edge based mechanisms can be

mitigated by using dynamic path capacity allocation algorithms that exploit knowledge of the traffic matrix.

- With increasing complexity in the way endpoints in a VPN interact, the importance of understanding the traffic matrix increases.
- Traffic matrix estimation is a dominant factor in determining the utilization gains. Deploying statistical admission control techniques might not be worth the effort if the traffic matrix is not incorporated.

Thus it is important to estimate the traffic matrix for VPNs. In the absence of signaling-based admission control mechanisms it is advisable to build a dynamic path allocation architecture as described here. Adopting such an approach, the performance gap between signaling and non-signaling mechanisms reduces considerably. In conclusion, our results help a designer choose the right pieces to build a provisioning strategy that yields higher resource utilization gains.

#### ACKNOWLEDGMENT

This work was supported in part by the DARPA grant F30602-00-2-0537 and AT&T.

#### REFERENCES

- [1] R. Boorstyn, A. Burchard, J. Liebeherr, and C. Oottamakorn, "Statistical service assurances for traffic scheduling algorithms," *IEEE J. Select. Areas Commun.*, vol. 18, no. 12, pp. 2651–2664, Dec. 2000.
- [2] CAIDA, "MapNet raw source data files." [Online]. Available: <http://www.caida.org/tools/visualization/mapnet/Data/>
- [3] N. Duffield, P. Goyal, A. Greenberg, P. Mishra, K. Ramakrishnan, and J. van der Merive, "Resource management with hoses: point-to-cloud services for virtual private networks," *IEEE/ACM Trans. Networking*, vol. 10, no. 5, pp. 679–692, Oct. 2002.
- [4] A. Elwalid, D. Mitra, and R. Wentworth, "A new approach for allocating buffers and bandwidth to heterogeneous, regulated traffic in an ATM node," *IEEE J. Select. Areas Commun.*, vol. 13, no. 6, pp. 1115–1127, Aug. 1995.
- [5] V. Firoiu, J.-Y. Le Boudec, D. Towsley, and Z. Zhang, "Theories and models for internet quality of service," in *Proc. of the IEEE*, vol. 90, no. 9, Sept. 2002, pp. 1565–1591.
- [6] A. Gupta, J. M. Kleinberg, A. Kumar, R. Rastogi, and B. Yener, "Provisioning a virtual private network: a network design problem for multicommodity flow," in *ACM Symposium on Theory of Computing*, 2001, pp. 389–398. [Online]. Available: <http://citeseer.nj.nec.com/article/gupta01provisioning.html>
- [7] G. Italiano, R. Rastogi, and B. Yener, "Restoration algorithms for virtual private networks in the hose model," in *Proc. IEEE INFOCOM 2002*, vol. 1, 2002, pp. 131–139.
- [8] E. Knightly, "Enforceable quality of service guarantees for bursty traffic streams," in *Proc. IEEE INFOCOM'98*, vol. 2, 1998, pp. 635–642.
- [9] E. Knightly and N. Shroff, "Admission control for statistical QoS: theory and practice," *IEEE Network*, vol. 13, no. 2, pp. 20–29, Mar. 1999.
- [10] A. Kumar, R. Rastogi, A. Silberschatz, and B. Yener, "Algorithms for provisioning virtual private networks in the hose model," in *Proc. of ACM SIGCOMM 2001*, 2001, pp. 135–146.
- [11] S. Raghunath, K. Chandrayana, and S. Kalyanaraman, "Edge-based QoS provisioning for point-to-set assured services," in *Proc. of ICC 2002*, vol. 2, Apr. 2002, pp. 1128–1134.
- [12] S. Raghunath and S. Kalyanaraman, "Statistical Point-to-Set edge-based quality of service provisioning," in *Proc. of QoS'03*, Springer Verlag LNCS 2811, vol. 2, Oct. 2003, pp. 132–141.
- [13] S. Raghunath, K. Ramakrishnan, S. Kalyanaraman, and C. Chase, "Measurement based characterization and provisioning of IP VPNs," in *IMC 2004*, Oct. 2004. [Online]. Available: [http://networks.ecse.rpi.edu/~rsatish/vpn\\_tm.ps](http://networks.ecse.rpi.edu/~rsatish/vpn_tm.ps)
- [14] M. Reisslein, K. Ross, and S. Rajagopal, "A framework for guaranteeing statistical QoS," *IEEE/ACM Trans. Networking*, vol. 10, no. 1, pp. 27–42, Feb. 2002.
- [15] Y. Zhang, M. Roughan, N. Duffield, and A. Greenberg, "Fast accurate computation of large-scale IP traffic matrices from link loads," in *Proc. of ACM SIGMETRICS 2003*, 2003, pp. 206–217.

#### APPENDIX I

##### STATISTICAL ADMISSION CONTROL TEST

We briefly present the statistical admission test reported in [14]. The worst case loss probability for a flow with peak rate  $\pi_i$  at a multiplexer with capacity  $C$  is given by:

$$\frac{1}{Cs^{*2} \sqrt{2\pi\mu_U''(s^*)}} e^{-s^*(C-\pi_i)+\mu_U(s^*)} \quad (1)$$

where  $s^*$  is the unique solution to  $\mu_U'(s^*) = C - \pi_i$  and for the set of flows  $I$  incident at the multiplexer,  $\mu_U$  is defined as:

$$\begin{aligned} \mu_U(s) &= \sum_{j \in I - \{i\}} \mu_{U_j}(s) \\ \mu_{U_j}(s) &:= \log E[e^{sU_j}] \\ U_j &= \begin{cases} \pi_j & \text{with probability } \frac{\rho_j}{\pi_j} \\ 0 & \text{with probability } 1 - \frac{\rho_j}{\pi_j} \end{cases} \end{aligned}$$

#### APPENDIX II

##### DERIVING PER-DESTINATION STATISTICS

In §III-D we discussed the specification of traffic matrix and mentioned that mean and variance of per-destination traffic fraction can be used to deduce its dual leaky-bucket description. Here we briefly present a method of deducing the per-destination dual leaky-bucket description. We are given that the aggregate  $A$  conforms to  $(\pi, \rho, \sigma)$  and that  $(m_j, v_j)$  are the mean and variance values of the fraction of traffic directed toward destination  $j$ . If the long-term average of the aggregate is given by  $\rho$ , its variance is bounded by  $\pi\rho - \rho^2$  (see e.g., [8], [12]). Thus we can obtain the mean and variance of  $A_j$  as follows:

$$\begin{aligned} E\{A_j\} &= E\{p_j A\} \\ &= m_j \rho \\ \text{Var}\{A_j\} &= E\{A_j^2\} - (E\{A_j\})^2 \\ &\leq m_j \rho \left( \pi \left( \frac{v_j}{m_j} + m_j \right) - m_j \rho \right) \end{aligned} \quad (2)$$

Thus a dual leaky-bucket specified as:  $\pi_j = \pi \left( \frac{v_j}{m_j} + m_j \right)$ ,  $\sigma_j = \sigma$  and  $\rho_j = m_j \rho$  represents the per-destination aggregate whose characteristics are described by  $p_j$  (the random variable denoting fraction of traffic toward  $j$ ).