# Statistical Point-to-Set Edge-Based Quality of Service Provisioning

Satish Raghunath, Shivkumar Kalyanaraman
Department of ECSE, Rensselaer Polytechnic Institute
Troy, NY 12180

***Abstract*— In this paper we propose an edge-based quality of service architecture aimed at site-to-site private networks over the Internet. We extend the traditional point-to-point service model to a point-to-set service model, assuming a finite, bounded set of destination sites. Instead of provisioning point-to-point links between a source and its set of destinations, a point-to-set service allows the user to have an allocated bandwidth, which could be flexibly assigned to traffic going toward any destination within the set. The proposed point-to-set service provides low loss rates and *flexibility* to users while allowing providers to obtain multiplexing gains by employing a probabilistic admission control test.**

**The model is demonstrated to be parsimonious in parameters and completely implemented at the edge of the network. We provide an intuitive measure to quantify the *flexibility* of a point-to-set service and demonstrate its utility in deciding the trade-off between low loss rates and high multiplexing gains. Simulation results are presented to demonstrate the merits of the proposed architecture in terms of loss and delay characteristics seen by the user and multiplexing gains for the provider.**

## I. INTRODUCTION

The best-effort traffic in Internet is inherently of the point-to-anywhere nature, i.e., sources direct packets to any possible destination. In contrast, traditional quality-of-service (QoS) models set up premium services on a *point-to-point* basis (eg: virtual leased lines, frame-relay, ATM services, int-serv [3] etc). Recently, with the advent of IP differentiated services [1], [6], [27] there has been interest in expanding the *spatial granularity* of QoS models. Clark and Fang [6] proposed that a pool of "assured" service tokens could be allocated to a user or site with the flexibility to employ the tokens toward any arbitrary destination. While such a *"point-to-anywhere"* assured service model is very appealing to users, the large spatial granularity of the service makes efficient admission control and provisioning virtually impossible [27]. We consider a subset of this problem by examining assurances to a fixed set of destinations and introducing limits on offered load toward each destination. We provide a novel solution wherein a user has considerable *flexibility* in apportioning the allocated bandwidth among the destinations in the pre-defined set and the provider also sees multiplexing gains. Unlike previous solutions (e.g., the Hose Model [8]) our model does not require bandwidth demand estimation or reservation and hence does not need signaling.

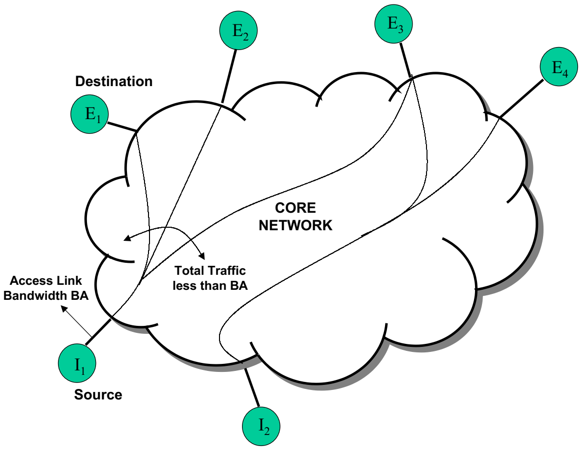

Fig. 1. The Point-to-Set Concept: The total offered traffic due to $I_1$ is limited by its access link bandwidth. Provider gains most if the reserved bandwidth is less than or equal to this quantity.

### A. The Point-to-Set concept

Consider a private network of sites $I_1$, $I_2$, $E_1$, $E_2$, $E_3$ and $E_4$ as shown in Fig. 1. The aggregate traffic from $I_1$ (called the *point*) toward $E_1$, $E_2$ or $E_3$ (called the *set*) is bounded by the capacity of the access link (say, "peak"). Given the point-to-point allocation model, site $I_1$ would require a link with capacity equal to "peak", to each destination in the *set* for an assured service toward the sites in the set. As such, the total purchased capacity from the provider (which is three times the peak here) exceeds the access link capacity leading to wastage of

---

Supported in part by NSF contract ANI9806660, NSF contract ANI9819112 and grants from Intel Corp. and Nortel Networks Limited

resources. We propose a *point-to-set service* wherein a customer buys a bandwidth *less than or equal* to his peak requirement (or a given total bandwidth), but is assured that his traffic needs to any destination in the set are met with a *probability* close to 1. In other words, the user buys bandwidth to a set of destinations, instead of purchasing point-to-point links to the destinations and retains the freedom of deciding the fraction of bandwidth allocated to a specific destination. Thus there is a cost saving in that the point-to-point links need not be leased from $I_1$ to each member of the set. For the provider, the paths connecting edge $I_1$ to the set $\{E_1, E_2, E_3\}$ can be multiplexed with other contracts by exploiting statistical properties of the traffic.

### *B. An Ideal Point-to-Set Service*

Before trying to build an architecture to realize the point-to-set concept, it is useful to consider the ideal implementation. A user would want to be assured a bandwidth equal to the peak requirement toward *any* destination. A more restricted version of the ideal case is where the set of destinations is finite and the user still has the assured bandwidth toward any node in this finite set.

Consider how a provider would implement this ideal service. An efficient provisioning strategy would reserve a network-wide total bandwidth less than or equal to the peak requirement of the customer. However, the user can offer traffic at this peak rate toward any destination. Since the user does not specify the exact load toward a given destination, the provider needs to accurately predict demand to avoid over-provisioning.

In previous work [22] we have investigated such an ideal strategy. In practice, such dynamic tracking and provisioning schemes are hard to implement due to complex and time-varying statistical characteristics of Internet traffic. The intuitive appeal of a point-to-set service is in the fact that it has the potential to provide inexpensive and *flexible* services to the customer while allowing statistical multiplexing gains to the provider. The important questions to be answered then are:

- Are there quantifiable benefits that make deployment of point-to-set services an attractive option?
- What are the simplifications to the ideal service that will render the architecture practical and realizable?
- Is it possible to build such a service with a minimal edge-based approach?

In the following sections, we build a point-to-set service with some simplifications to the ideal model so that it can be implemented at the edge of the network with just simple shaping components.

### *C. Building A Deployable Model*

In the ideal point-to-set model, the onus of gathering information regarding user traffic is completely on the provider. In order to simplify the model from the perspective of making the network simple, the user could be required to conform to a certain traffic profile.

The resource wastage in a point-to-point allocation model is due to over-provisioning caused by lack of knowledge regarding the fraction of total load offered toward a destination. The solution to this could be in assuming something about the per-destination load. At one extreme would be the choice of assuming that a fixed fraction of the total traffic is offered toward each destination; this is no better than the point-to-point model in terms of either flexibility or multiplexing gain. In order to allow for the dynamic nature of traffic we could strike a *middle ground* between the two extremes of assuming all or nothing about the per-destination traffic. We could assume that the fraction of traffic toward a given destination is random, but has a given mean and variance $(m, v)$.

This approach would allow the traffic fraction toward a destination to vary within the limits specified by $(m, v)$. Further, knowing the leaky-bucket parameters shaping the total traffic, one could compute bounds on the probability of observing a particular load toward a given destination. We employ this approach to demonstrate that a simple probabilistic admission control scheme can be derived in terms of the mean and variance parameters of per-destination traffic fraction and the leaky bucket parameters of the total traffic. We then show that these $(m, v)$ parameters can be enforced using simple deterministic shaper elements.

Note that a higher variance for the per-destination fraction implies greater freedom to the user as regards to bandwidth usage. This is in essence *higher flexibility*. Exploiting this intuition, we define flexibility as an upper bound on the variance to mean ratio of per-destination traffic fraction. We demonstrate via simulations that this definition satisfies all the intuitive requirements of such a measure. We then explore the role of flexibility in the trade-off between loss rate and multiplexing gains.

The contributions of this paper are thus as follows: a) A novel architecture for statistical edge-based bandwidth provisioning toward a set of destinations; b) A simple means to capture and enforce the per-destination traffic statistics; c) A probabilistic admission control test that allows a flexible service for the user with simultaneous multiplexing gains to the provider; and d) A solution that can be deployed at the edge nodes of the provider network without altering the core.

## II. RELATED WORK

In this section we review literature concerning analytical frameworks for statistical quality of service and architectural proposals for the Internet. We relate our work to previous proposals and demonstrate the novelty in our contributions.

### A. Network Architectures for QoS

Clark et al [6] introduced the idea of going beyond point-to-point services and providing flexibility to users, while at the same time allowing multiplexing gains for the provider.

LIRA [27] considers the problem of large spatial granularity, where QoS assurances are for a large set of destinations (possibly unlimited). By employing enhancements to routing protocols the authors provide a way to achieve per-packet admission control so that a user can employ the allocated bandwidth toward any destination. Consequently, LIRA faces scalability issues when there are large number of destinations. In the present paper, we do not require any changes to the core network or the routing protocols. Further, we consider admission control on aggregates and assurances toward a finite set of destinations. Hence our proposal is not affected by most of the scalability issues mentioned above.

Duffield et al [8] propose a framework for Virtual Private Network (VPN) resource management and introduce the idea of a "hose" as a resizeable access link for a VPN node. They attempt to solve a part of the problem tackled here, namely, that of going beyond the point-to-point allocation model and do not treat the problem of admission control. The hose is intended to provide bandwidth toward the set of destinations and is implemented by the provider using a reservation tree structure [15]. They suggest resizing the hose using online prediction of traffic characteristics to obtain further multiplexing gains. The hose model exploits multiplexing gain while providing a single logical interface between the customer and provider and eliminates the need to specify a traffic matrix. But the Hose Model requires solving a complex optimization problem to find a reservation tree structure for the purposes of provisioning. Further, the reservation is set up using signaling and per-link admission control. Since the hose is resizeable, every change in allocated bandwidth is accompanied by per-link admission control. The resizing itself depends on a Gaussian bandwidth demand predictor. Thus the performance of the model is dependent on traffic statistics being amenable to the predictor assumptions.

| **Attribute** | **Customer Pipe** | **Hose Model** | **Point-to-Set** |
|---|---|---|---|
| Deployment | Point-to-Point links for each source dest. pair | A single *hose* from customer; network-wide reservation | Single interface with customer; Fully Edge-based |
| Bandwidth Reservation | Static - Whole link is reserved | Dynamic dep. on demand | No reservation |
| Signaling | None | Required to update reservations | Not Required |
| Traffic Matrix | Need info about every source-dest pair | Not reqd. | Need mean variance of traff. fraction |
| Traffic Statistics | Provision for peak traffic | Gaussian predictor to track demand | Does not need any online traff. stats. |
| Algorithmic Complexity | None | Complex provisioning algorithm | No b/w reservation |
| Admission Control | Deterministic, one-time | Need per-link per-hose computation for every change in reservation | Edge-based, statistical, one-time computation |
| Multiplexing gains | None | Statistical gains due to hose-level aggregation | High gains due to statistical admission ctl |

TABLE I

COMPARISON OF MODELS FOR HANDLING QOS WITH INCREASED SPATIAL GRANULARITY.

In contrast, the Point-to-Set model adopts a very different approach and does not rely on demand estimation or signaling based reservations. The Point-to-Set model depends instead on a probabilistic admission control regime to exploit multiplexing gain. As a consequence, the complexity due to reservation tree structures and per-link admission control are eliminated. The simplification comes at the cost of additional traffic information in the form of mean and variance of the per-destination traffic fraction. This information can be obtained in a characterization phase as described in [8]. Table I summarizes the preceding discussion.

### B. Statistical Quality of Service

Statistical admission control schemes have been shown to achieve much higher utilization than the deterministic counterparts [14]. Such a scheme arrives at a decision by verifying that QoS metrics of interest to the network are not violating the specified limits. With input being characterized statistically, the task of computing end-to-end statistical bounds on QoS metrics becomes complex due to the correlation amongst flows

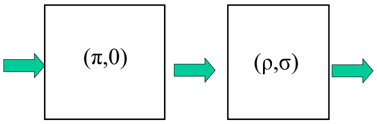

Fig. 2. A dual-leaky-bucket regulator has two shapers in series.

that exit a multiplexer. To make the problem tractable one can either avoid correlation amongst flows (e.g., using bufferless multiplexers [24], jitter control at each node [10], [17], [29]) or resort to approximate analysis (e.g., busy period analysis [16], large deviation methods [28]).

With the assumption of independence amongst flows, statistical admission control conditions have been evaluated for sources bound by certain probabilistic envelopes (e.g., Rate-Variance envelopes [12], Effective envelopes [4]). These envelopes are easily computed for mutually independent flows given the underlying deterministic envelope (e.g. leaky-bucket shaper parameters).

In order to build a probabilistic admission control mechanism for the point-to-set architecture, we adopt a dual-leaky-bucket regulated source characterization (similar to [25]) and relate the parameters to statistical characteristics. Unlike existing work, we obtain bounds on per-path traffic statistics at the edge of the network exploiting the point-to-set model. We then employ the per-path information to evaluate the admission control criterion.

## III. THE POINT-TO-SET ARCHITECTURE

We first outline the assumptions and notations for the rest of the paper (§III-A). After an overview of the architecture (§III-B) the components are defined (§III-C).

### A. Notations and Assumptions

Table II provides a brief description of the symbols that are used in the succeeding sections. In the following sections, a “user” refers to a customer network offering traffic. A “flow” is a traffic aggregate emanating from a network. The user traffic is assumed to be shaped by a dual-leaky-bucket regulator of the form $(\pi, \rho, \sigma)$ (Fig. 2). Thus, the cumulative offered traffic $A(t)$ in time $t$ always satisfies $\{A(t) \leq \pi t, A(t) \leq \rho t + \sigma\}$. This is equivalent to having $(\rho, \sigma)$ and $(\pi, 0)$ leaky-bucket shapers in series. A QoS commitment to the user is termed as a contract (defined in §III-C).

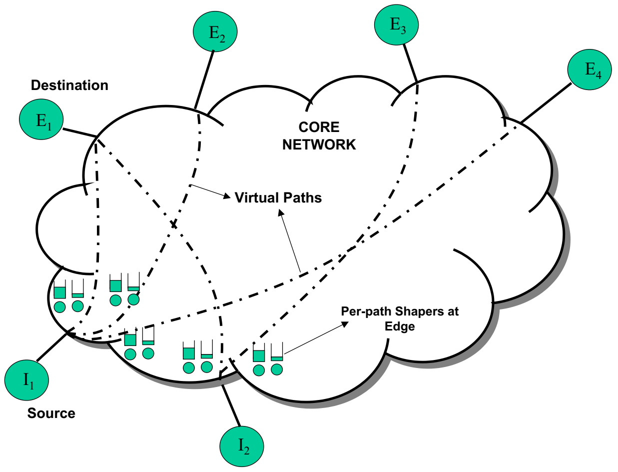

Fig. 3. The Point-to-Set Architectural Model consists of dual-leaky-bucket regulators per-path for a source network offering traffic. In the figure, traffic from network $I_1$ is directed toward $E_1$,$E_2$ and $E_4$. Each of these virtual paths is regulated at the ingress.

The admission control module is assumed to know the paths connecting ingresses and egresses. In this paper we shall assume that all packets destined to a particular egress from a given ingress use the same path (i.e., at a given time there is a single path connecting an ingress to an egress). Hence each “path” is uniquely associated with an ingress-egress pair. Routes are assumed to remain stable.

### B. Overview

The point-to-set architecture is depicted in Fig. 3. Each user network that enters into a contract with the provider is assumed to specify the set of destinations and the mean and variance of per-destination traffic fraction. This fraction is enforced via dual-leaky buckets as shown in Fig. 3.

The admission control module is a central entity to the provider network that knows the paths connecting the provider edge nodes. A contract requires bandwidth provisioning toward every destination in its set. Consequently, a new contract can be admitted only if the bandwidth requirement can be accommodated along each *path* connecting the ingress node to a destination. Thus while deciding to admit a contract, the module checks whether adding this contract would cause input rate to exceed the “path capacity”. We define path capacity in the next section. Here we provide an intuitive description of the concept.

A given path from an ingress to an egress may share one or more physical links with other paths in the network. Hence its capacity is not the same as that of the physical links constituting the path. For admission

| Symbol | Meaning |
|---|---|
| $\pi_j$,$\rho_j$, $\sigma_j$ | Peak rate, Avg rate, Bucket for user $j$ |
| $\pi_{ij}$,$\rho_{ij}$, $\sigma_{ij}$ | Peak, Avg, Bucket for user $j$ toward dest $i$ |
| $p_{ij}$ | Random variable for traffic fraction toward $i$ for user $j$ |
| $m_{ij}$ | Mean of fraction of load toward dest $i$ for user $j$ |
| $v_{ij}$ | Variance of traffic fraction toward dest $i$ for user $j$ |
| $C_i$ | Capacity of path $i$ |
| $D_{max}$ | Max permissible delay at ingress |
| $\epsilon$ | Max allowed capacity violation probability |
| $Y_j$ | Random variable for total traffic (bps) due to user $j$ |
| $X_{ij}$ | Random variable for traffic (bps) toward node $i$ for user $j$ |
| $F$, $f$ | Flexibility |
| $\Gamma_j$ | Capacity of link $j$ |

TABLE II

TABLE OF NOTATIONS

control purposes, it is convenient to introduce a notion of a *virtual path* of fixed capacity connecting the ingress to the egress. In other words, a virtual path is a means to apportion link capacities among paths. As shown in Fig. 3 a virtual path between an ingress-egress pair appears as if it is dedicated to this pair.

Then the task of the admission control test is to verify that the probability that the input traffic exceeds the path capacity is less than a given threshold for every path affected by the new contract. The idea of a virtual path thus allows us to achieve admission control at the edge of the network.

### C. Definitions

We first define the path capacity and the contract specification.

*Definition 3.1: Consider a path defined by the sequence of links $\{M_j\}$. Let $\Gamma_j$ be the capacity of $M_j$. Let $n_j$ be the number of paths passing through $M_j$. Then, the capacity of the path, is defined as:* $C = \min_j \frac{\Gamma_j}{n_j}$.

The capacity of a path is a fixed constant once the topology and routing are fixed. We note that the above definition of path capacity is very simple and has some drawbacks. This definition accounts only for the number of paths sharing the bottleneck network links. A better definition would attempt to apportion the bottleneck bandwidth in proportions that would yield higher capacity to paths that are more “important” (e.g., have higher link bandwidths preceding the bottleneck). However, in succeeding sections we shall use this definition and treat such an improved algorithm in future work.

*Definition 3.2: A Contract for user network $j$ consists of the dual-leaky-bucket characterization of the total traffic given by $(\pi_j, \rho_j, \sigma_j)$, the finite set of destination nodes, $S_j$, the set of pairs $\{(m_{ij}, v_{ij}) \mid i \in S_j\}$ where $(m_{ij}, v_{ij})$ are the mean and variance of the random variable $p_{ij}$ indicating the fraction of total traffic toward $i$. So if total traffic (bits/second) is given by $Y_j$ and $X_{ij}$ indicates the traffic toward destination $i$, $X_{ij} = p_{ij}Y_j$, $p_{ij} \in (0, 1]$ and $\sum_i X_{ij} = Y_j$.*

### D. Admission Control Test

The key idea that we exploit here is that of getting an *a priori* estimate of the fraction of total traffic that is offered along a given path. Denote the total traffic (bits per second) offered by customer $j$ as $Y_j$. Let $X_{ij}$ denote the traffic due to customer $j$ on the path leading to the destination $i$. If $p_{ij}$ are fixed constants, the provider can provision the right amount of bandwidth toward each destination. This would be identical to a point-to-point service toward each of the destinations. A more interesting and realistic situation is when $p_{ij}$ are not fixed.

Thus we define $p_{ij}$ as a random variable with mean and variance $(m_{ij}, v_{ij})$. For simplicity, we assume $p_{ij}$ are independent of $Y_j$, i.e., the fraction of traffic toward a destination is independent of the total volume of traffic offered by the network. We now impose the constraint that $Y_j$ is policed to a peak rate $\pi_j$ and shaped by a leaky-bucket shaper $(\rho_j, \sigma_j)$.

Our goal is to reserve only as much bandwidth as a customer offers toward a destination. Thus our admission control strategy should consider the traffic that a customer might provide toward a particular destination and also attempt to exploit multiplexing gains. Using the random variables $X_{ij}(t)$ we could formulate an admission control condition as follows - admit a new contract if:

$$\forall i,\ Pr\{\sum_j X_{ij}(t) > C_i\} \ < \ \epsilon \tag{1}$$

$$\forall i,\ \sum_j m_{ij}\rho_j \ < \ C_i \tag{2}$$

Here $\epsilon < 1$ is a given constant. In essence, this condition is measuring the probability that the total offered load on a path ($\sum_j X_{ij}(t)$) exceeds the capacity of that path ($C_i$), given that the average rate of admitted contracts ($\sum_j m_{ij}\rho_j$) does not exceed the path capacity.

Observe that Equation (1) serves our objectives well. First, it reserves per-path bandwidth depending on the amount of traffic that the contract might offer. Second, it allows us to exploit *statistical* multiplexing gains

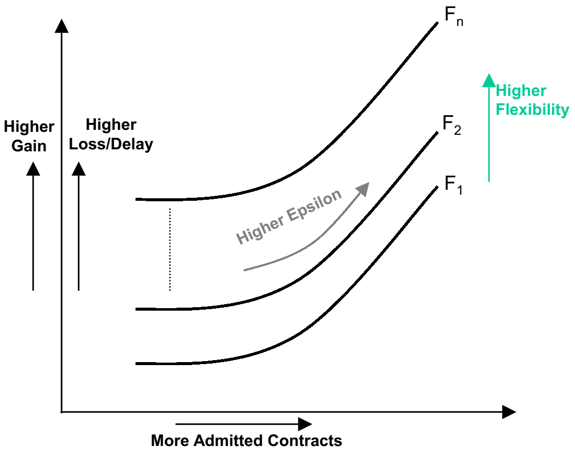

Fig. 4. Schematic showing the significance of Epsilon ($\epsilon$) and Flexibility. Higher flexibility requires lesser number of admitted contracts if loss rates and delays have to be maintained at the same level.

by not choosing peak provisioning. The parameter $\epsilon$ provides a control on how conservative the admission control gets. Lower the value of $\epsilon$ higher the reserved bandwidth and lower the multiplexing gain. Also note that the equation inherently captures the fact that a customer network might vary the fraction of traffic it sends along a given path. Comparing this strategy with a deterministic strategy readily points us to the gains in exploiting the varying nature of customer's traffic. A deterministic point-to-point service toward each destination would need a fixed $p_{ij}$ or would have to choose peak provisioning. In §V we quantify these gains using simulations.

The relation of $\epsilon$ to the loss rate experienced along a path is not so simple due to the distortions introduced in traffic characteristics by multiplexing at successive hops. We shall return to this aspect of loss rates in §IV-A.

### E. *Quantifying Flexibility*

An ideal Point-to-Set service provides an abstraction of a point-to-point link toward each destination for a contract. The source network has the *flexibility* to offer an arbitrary fraction of its total traffic toward any destination. In a realistic implementation, there will be a limit to the variability in the source network's traffic. To measure how close to ideal an implementation is, we could examine the flexibility it offers. We expect a measure of flexibility to satisfy these intuitive requirements:

- Higher the flexibility, greater is the freedom to the user in terms of load distribution with respect to the destinations.
- If loss rates are kept low, higher the flexibility, closer is the service to a point-to-point regime.

Define flexibility, $f$ so that:

$$\frac{\sqrt{v_{ij}}}{m_{ij}} \leq f \ \forall i, j$$

The definition implies that a higher value of $f$ allows for higher variance in per-path offered load. In order to attain lower loss rates and still allow for higher $f$ one would have to admit lesser number of contracts, i.e., employ a lower value of $\epsilon$. On the other hand, allowing higher variances for a given set of admitted contracts can lead to higher loss rates. Thus, the following can be stated as the properties of $f$:

- For the same multiplexing gain, higher the flexibility, higher the loss rates to the users.
- For the same loss rate, higher the flexibility lower the multiplexing gain.

These properties are captured in the schematic diagram in Fig. 4. As we go along one of the curves, we are holding flexibility constant while increasing the violation probability $\epsilon$ and hence increasing multiplexing gain and loss rate. If we move up vertically (increase flexibility) for the same number of admitted contracts, we again increase loss rates. These observations are verified through simulations in §V. The preceding discussion thus points to a trade-off between flexibility and loss rate.

## IV. EVALUATING THE ADMISSION CONTROL DECISION

In the following paragraphs we derive approximations that will help us evaluate the admission control test in Equation (1).

### A. *Per-Path Traffic Statistics*

In order to evaluate Equation (1), we would need the distribution of $X_{ij}$. An alternative approach would be to *bound* the distribution somehow, exploiting the fact that $Y_j$ was constrained by $(\pi_j, \rho_j, \sigma_j)$. We thus obtain an upper bound on the mean and variance of the process. To do this we employ a technique similar to [13] and observe that the extremal "on-off" source (Figure 5) has the maximum variance among all rate patterns that can be obtained given the $(\pi_j, \rho_j, \sigma_j)$ characterization, if mean is set at $\rho_j$ (Proposition 1). Our approach differs from that of [13] in having a bound independent of a specific interval or duration, and in considering the dual leaky-bucket shaped inputs specified by $(\pi_j, \rho_j, \sigma_j)$.

We then employ this upper bound on the variance of rate to obtain a bound on per-path traffic statistics

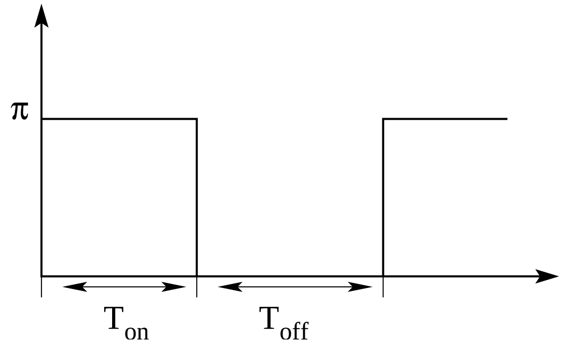

Fig. 5. The Extremal On-Off Source

(Proposition 2). Equipped with this result we consider an approximate evaluation of Equation (1) (Proposition 3).

We note that although the extremal on-off source has maximum variance it does not necessarily maximize buffer overflow probability [7], [11], [23]. The extremal source has been used in the past [9], [18], [25] with reference to bandwidth and buffer allocation. Here we employ the source owing to the fact that it leads to more conservative provisioning while easing analysis.

*Proposition 1: Consider a source shaped as* $(\pi_j, \rho_j, \sigma_j)$. *A transmission pattern with mean* $\rho_j$, *that maximizes the variance of rate is given by the periodic extremal on-off source, wherein the source transmits at the peak rate* $\pi_j$ *for a duration* $T_{on} = \frac{\sigma_j}{\pi_j - \rho_j}$ *and switches off for* $T_{off} = \frac{\sigma_j}{\rho_j}$.

*Proof:* Consider the density function $f_X(x)$ corresponding to the extremal source and its variance $v_X$:

$$\begin{aligned} f_X(x) &= \frac{T_{off}}{T_{on}+T_{off}}\delta(x) + \frac{T_{on}}{T_{on}+T_{off}}\delta(x-\pi_j) \\ v_X &= \pi_j^2 \frac{T_{on}}{T_{on}+T_{off}} - \rho_j^2 \\ &= \pi_j\rho_j - \rho_j^2 \end{aligned} \tag{3}$$

Let $f_Y(y)$ denote any other density function such that $Y$ is shaped according to $(\pi_j, \rho_j, \sigma_j)$ and has mean $\rho_j$. Compare its variance, $v_Y$ with that of $X$:

$$\begin{aligned} v_X - v_Y &= E\{X^2\} - E\{Y^2\} - (E\{X\})^2 + (E\{Y\})^2 \\ &= E\{X^2\} - E\{Y^2\} \\ &= \pi_j\rho_j - \int_0^{\pi_j} y^2 f_Y(y)dy \\ &= \int_0^{\pi_j} \pi_j y f_Y(y)dy - \int_0^{\pi_j} y^2 f_Y(y)dy \\ &= \int_0^{\pi_j} y(\pi_j - y) f_Y(y)dy \\ &\geq 0 \end{aligned}$$

■

With this proposition, we can now consider the first and second moments of the per-path traffic due to a contract, namely, $X_{ij}$. The statistical characteristics of traffic is altered by each hop of multiplexing. Multiplexing introduces correlation among flows and increases burstiness [2]. Although the mean remains the same, the variance of rate is higher at a node further along a path in the network. This has implications on provisioning buffers inside the network. We can evaluate Equation (1) to ensure that the mean of admitted traffic remains below the path capacity and the buffer requirement at the edge of the network is low. To achieve low loss rates, buffers inside the network have to be appropriately set. We first examine Equation (1) and treat buffer dimensioning in §IV-C.

*Proposition 2: If* $Y_j$, *the total traffic due to customer* $j$, *shaped by a dual leaky-bucket shaper* ($\pi_j$,$\rho_j$,$\sigma_j$) *has a mean* $\rho_j$ *and* $X_{ij}$ *is the fraction of* $Y_j$ *along path* $i$, *the mean and variance of* $X_{ij}$ *are given as follows.*

$$E\{X_{ij}\} = m_{ij}\rho_j \tag{4}$$

$$Var\{X_{ij}\} \leq m_{ij}\rho_j(\pi_j(\frac{v_{ij}}{m_{ij}} + m_{ij}) - m_{ij}\rho_j) \tag{5}$$

*Proof:*

$$\begin{aligned} E\{X_{ij}\} &= E\{p_{ij}Y_j\} \\ &= E\{p_{ij}\}E\{Y_j\} \\ &= m_{ij}\rho_j \\ Var\{X_{ij}\} &= E\{X_{ij}^2\} - (E\{X_{ij}\})^2 \\ &= E\{p_{ij}^2\}E\{Y_j^2\} - m_{ij}^2\rho_j^2 \\ &\leq (v_{ij} + m_{ij}^2)\pi_j\rho_j - m_{ij}^2\rho_j^2 \\ &= m_{ij}\rho_j(\pi_j(\frac{v_{ij}}{m_{ij}} + m_{ij}) - m_{ij}\rho_j) \end{aligned} \tag{6}$$

■

Observing that for each path, the statistical characteristics of the traffic offered by a given customer is independent of those of others at the *edge of the network* we now propose an approximation.

*Proposition 3: Define the Gaussian random variable* $Z_i$ *with mean* $m_{Z_i} = \sum_j m_{ij}\rho_j$ *and variance* $v_{Z_i} = \sum_j m_{ij}\rho_j(\pi_j(\frac{v_{ij}}{m_{ij}} + m_{ij}) - m_{ij}\rho_j)$. *Then for sufficiently large number of admitted customers, we have the following approximation.*

$$Pr\{\sum_j X_{ij} > C_i\} \leq Pr\{Z_i > C_i\}$$

$$\approx \frac{1}{\sqrt{2\pi}} exp\left(-\frac{(C_i - m_{Z_i})^2}{2v_{Z_i}}\right) \tag{7}$$

*Proof:* Since $X_{ij}, \forall j$ are independent, $Var\{\sum_j X_{ij}\}$ is given by $\sum_j Var\{X_{ij}\}$ which is less than $v_{Z_i}$ as defined in Proposition 3. Note that $Var\{X_{ij}\}, \forall j$ can be assumed to be small compared to

$v_{Z_i}$ for sufficiently large number of customers. Then we can invoke the Central Limit Theorem to approximate Equation (1) by the Gaussian complementary cumulative probability as specified above in Equation (7). ∎

### B. Enforcing the per-path limits

Once a contract is admitted, the provider needs to ensure that the offered traffic adheres to the per-path mean and variance restrictions. Thus we need shaping elements that will enforce the terms of the contract, viz., the per-path mean and variance specified by $(m_{ij}, v_{ij})$. As demonstrated by the following proposition, it is straightforward to derive the dual leaky bucket shaper $(\pi_{ij}, \sigma_{ij}, \rho_{ij})$ for the path $i$ in terms of $(m_{ij}, v_{ij})$ and $(\pi_j, \rho_j)$.

*Proposition 4: Define the dual leaky bucket shaper* $(\pi_{ij}, \sigma_{ij}, \rho_{ij})$ *such that:*

$$\pi_{ij} = \pi_j\left(\frac{v_{ij}}{m_{ij}} + m_{ij}\right) \tag{8}$$

$$\sigma_{ij} = \sigma_j \tag{9}$$

$$\rho_{ij} = m_{ij}\rho_j \tag{10}$$

*This dual leaky bucket shaper ensures that the per-path traffic fraction with mean* $m_{ij}\rho_j$ *has variance less than* $m_{ij}\rho_j(\pi_j(\frac{v_{ij}}{m_{ij}} + m_{ij}) - m_{ij}\rho_j)$

*Proof:* Denote the variance of the output process of this shaper by $v$. From Equation (3) we see that

$$\begin{aligned} v &\leq \rho_{ij}(\pi_{ij} - \rho_{ij}) \\ &= m_{ij}\rho_j\left(\pi_j\left(\frac{v_{ij}}{m_{ij}} + m_{ij}\right) - \rho_{ij}\right) \end{aligned}$$

∎

With the above proposition, we now have the ability to implement the model with simple shaping elements.

### C. Buffer Dimensioning

In order to decide the size of buffers at each hop, we can either set a limit on the maximum tolerable per-hop delay or constrain the maximum burstiness of the input traffic at each node.

Let the maximum tolerable per-hop delay be $D_{max}$. The corresponding buffer size at a multiplexer serving at rate $C$ would be given by $C \times D_{max}$. While this strategy is simple and limits the maximum delay incurred, it can result in higher loss rates owing to increased burstiness inside the network.

The alternative of limiting the input burstiness at a given node inside the network is slightly more involved. We first note that the worst-case burstiness of a flow at the exit of a node increases in proportion to the sum of the burst characterizations of other flows being served by the same node. To limit the burstiness of a flow incident at a given node, we must limit the increase in burstiness due to every previous hop through which this flow passed. We do this by limiting the maximum increase in burstiness at the ingress.

Consider a multiplexer $M$. Let $P$ denote the set of multiplexers feeding traffic to $M$ and $L$ denote the set of incident flows at $M$. Let $D_i^{max}$ denote the maximum tolerable delay at multiplexer $i$. We can set the buffer size at a multiplexer $i$ to $\sum_{l \in L} \sigma_l$ or a quantity that upper bounds it, as given below.

$$\begin{aligned} \sum_{l \in L} \sigma_l &= \sum_{p \in P} \sum_{l \in p} \sigma_l \\ &\leq \sum_{p \in P} D_p^{max} C_p = D_M^{max} \end{aligned} \tag{11}$$

If we set $D_{max}$ to be the maximum tolerable delay at every ingress, we can recursively compute the bound given in Equation (11) for a specific topology.

By using the second strategy, we observe that higher buffers are allocated at a multiplexer further along a path. Thus loss rates are reduced as compared to the first strategy. However, the trade-off is the assumption that the paths between every ingress and egress be known. Since this assumption is required to compute path capacities for admission control, it does not increase the complexity of the service.

## V. PERFORMANCE EVALUATION

In this section we verify working of the model using extensive simulations. In the succeeding sections the performance evaluation is performed with the following objectives:

- To verify the superiority of the probabilistic admission control condition in terms number of admitted contracts in comparison to point-to-point allocation model (§V-B).
- To validate the intuition behind the definition of flexibility (§V-C).
- To examine the role of $\epsilon$ as a “control knob” on how conservative the provisioning gets, i.e., lower $\epsilon$ should give us lower losses and delays with lower multiplexing gain and vice-versa (§V-D).
- To understand the aspects of utilization (viz., average and maximum) affected by varying $\epsilon$ and flexibility (§V-E).
- To study the effect of bias in offered load toward a few destinations in the destination set on multiplexing gains (§V-F).

We begin by detailing the method used to setup the simulations.

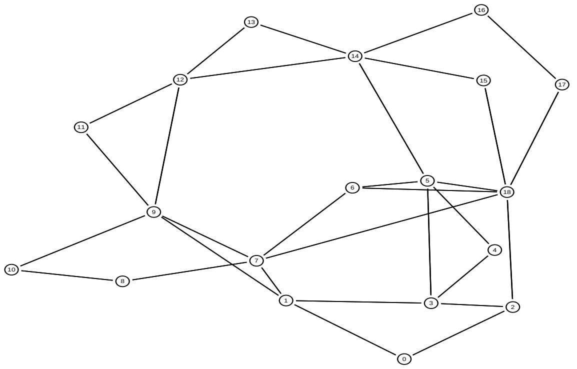

Fig. 6. The MCI topology used in simulations. Link capacities were set to 10 $Mbps$ and propagation delay was set to 10 $ms$

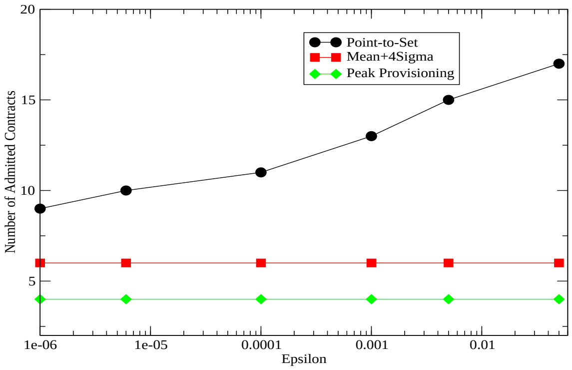

Fig. 7. Number of Admitted Contracts increases with increasing epsilon. The probabilistic admission control beats both $mean + 4 * sigma$ and peak provisioning

### A. *Methodology*

In order to evaluate the scheme, we employ Auckland IV traffic traces [19] [20] with the MCI backbone topology in NS-2 (shown in Fig. 6) simulation environment [21]. The Auckland data trace is a good fit for this evaluation due to a couple of reasons. First, it is a record of traffic at an access link and hence corresponds to the traffic generated by a network toward a set of destinations. Second, we need the generated traffic to be bursty; synthetically generated traffic might be too bursty (or may not be bursty enough). Using a real trace saves us from making decisions on the extent of traffic burstiness.

Each simulation consists of two phases - an admission control phase and a traffic generation phase. The simulation is started with a set of values for flexibility, $\epsilon$ and $D_{max}$ and is provided with randomly generated contracts. As examined earlier the contract consists of the set of destinations, the dual-leaky-bucket parameters for the total traffic and per-destination mean and variance for the traffic fraction. The contracts are admitted one after another until the admission control test fails. Then the traffic generation phase starts where the network performance metrics are measured for the admitted contracts.

To generate a contract randomly the following procedure was followed. For a destination set with 4 nodes, three uniform random numbers, $r_i, i = 2 \ldots 4$ are generated in the range $[min, max]$. Then setting $r_1 = 1$ and $\sum_i r_i m_{1j} = 1$ we obtain $m_{ij} = r_i m_{1j}$. For a given flexibility $v_{ij}$ can then be computed. The total traffic is then apportioned according to a Normal random variable with mean and variance $(m_{ij}, v_{ij})$ with negatives mapped to a small positive fraction.

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. This aspect helps us gauge the effect of bias on utilization - a higher bias can be expected to cause lower utilization.

In the simulations, the dual-leaky-bucket regulator parameters for all contracts was set at $(0.75\ Mbps,\ 0.5\ Mbps,\ 100\ kb)$. The link capacities were set to 10 $Mbps$ and their delay was chosen to be 10 $ms$. In the succeeding sections, each point in a graph indicating a simulation result, is the average of 10 simulation runs.

### B. *Comparing with the Point-to-Point Model*

The motivation for deploying point-to-set services is in the fact that there are multiplexing gains for the provider. In order to examine this aspect for the MCI topology, we compare the number of admitted contracts in a point-to-set service to that in a point-to-point service.

A point-to-point service provisions links at peak rate toward each of the destinations in the set. In addition to this, we introduce a model where the provider reserves $m_{ij} + 4\sqrt{v_{ij}}$ instead of doing a probabilistic admission control. Although this scheme is deterministic, it exploits the additional information regarding per-destination traffic fraction. Fig. 7 shows the number of admitted contracts under these three schemes and clearly a probabilistic scheme performs much better.

### C. *$\epsilon$ and Flexibility as Control Knobs*

While introducing $\epsilon$ and flexibility as parameters we noted that they provide a handle on the trade-off between loss rates and multiplexing gains. Here we seek to support those observations. In order to do this, we assign a fixed value to epsilon and study the number of contracts that can be admitted for various values of flexibility.

Intuitively, if a higher value of flexibility is allowed, the variance of per-destination traffic fraction can be

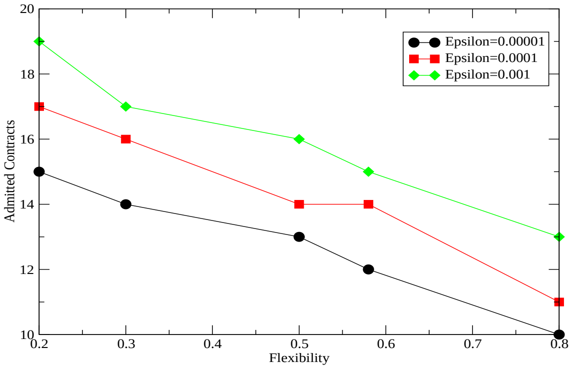

Fig. 8. For a fixed violation probability ($\epsilon$), higher flexibility implies lesser number of admitted contracts.

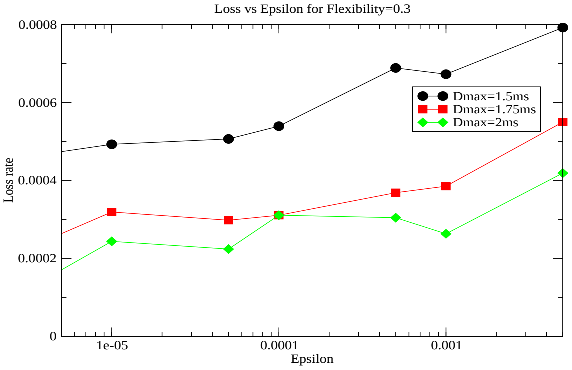

Fig. 9. Losses increase with more admitted contracts (increasing $\epsilon$) and lower buffer sizes (decreasing $D_{max}$).

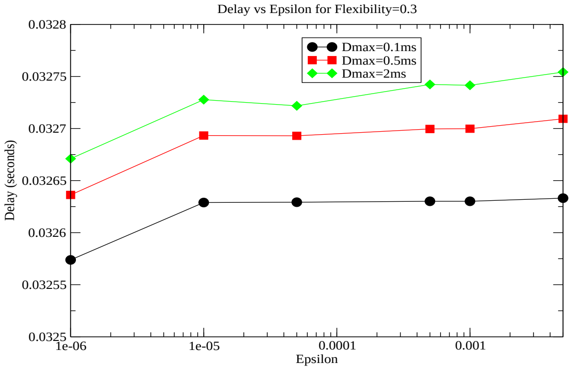

Fig. 10. Higher buffer sizes ($D_{max}$) and more number of admitted contracts imply higher average end-to-end delays.

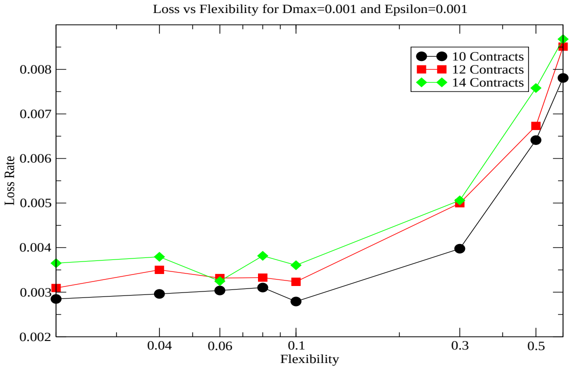

Fig. 11. Maintaining losses at roughly the same level with increase in flexibility requires admitting lesser number of contracts.

higher. This points to the fact that for the same $\epsilon$ lower number of contracts will be admitted. This is observed in Fig. 8. As before, with higher values of epsilon, the number of admitted contracts is higher.

This agrees with our initial description of flexibility as a measure in Fig. 4. We examine the effect of higher $\epsilon$ and flexibility on loss rates and delays in the succeeding sections.

### *D. Effect of Parameters on Loss and Delay*

In the preceding sections we observed that the $\epsilon$ and flexibility can be used to increase or decrease the number of admitted contracts. This implies that these two parameters serve as a handle on how conservative the provisioning gets. For the provider, these parameters present a trade-off between multiplexing gains and loss rates. For the user, flexibility offers a trade-off between freedom with respect to per-destination load variation and cost of the service.

Fig. 9 demonstrates the variation of loss rates with $\epsilon$ for different buffer sizes. Recall that $D_{max}$ decides the maximum permissible delay at the ingress and hence the buffer sizes (§IV-C). With higher buffers, as expected, loss rates are lower. In the present simulations, a $D_{max}$ of $2ms$ is shown to reduce the loss rates considerably. Loss rates consistently increase with higher $\epsilon$ since there is higher multiplexing.

The reduction in losses in Fig. 9 with higher buffer sizes comes at the cost of increased delays. As seen in Fig. 10, the average end-to-end delay experienced increases with higher $\epsilon$ and $D_{max}$. Thus the setting of $D_{max}$ and $\epsilon$ allows the provider to trade-off loss and delay with multiplexing gain.

Fixing $\epsilon$ and flexibility sets an upper-bound on the number of admissible contracts. If the provider chooses to allow for a higher flexibility he must admit lesser contracts to maintain the probability of capacity violation at $\epsilon$. Hence higher flexibility for the same number of admitted contracts comes at the cost of lower multiplexing gain. In addition, the higher variance in per-destination load leads to higher losses and delay. In Fig. 11 and Fig. 12 these aspects are demonstrated. To provide the same loss and delay characteristics at higher flexibilities, the number of contracts must reduce.

### *E. Utilization*

A higher $\epsilon$ allows admission of more number of contracts and hence allows for increasing the average

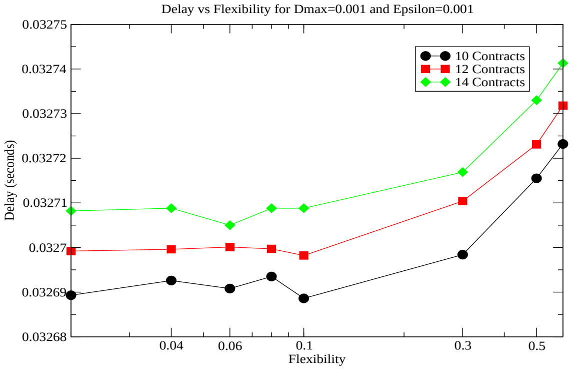

Fig. 12. Keeping delays low with increasing flexibility requires admitting lesser number of contracts.

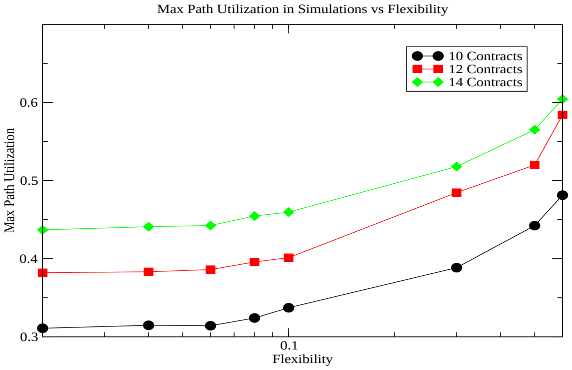

Fig. 13. Although the average utilization remains the same, increasing flexibility allows the maximum utilization levels to be higher. Increasing $\epsilon$ provides an additional dimension in which to raise maximum utilization levels.

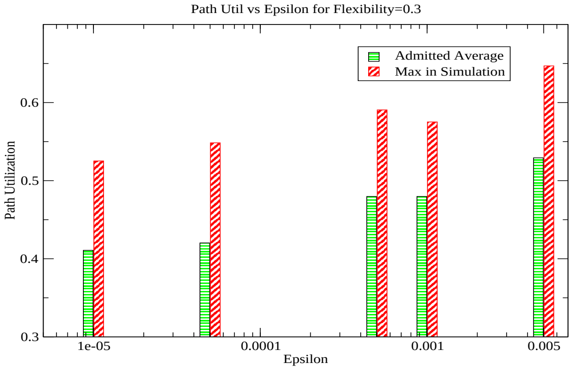

Fig. 14. Average Path Utilization increases with increasing $\epsilon$.

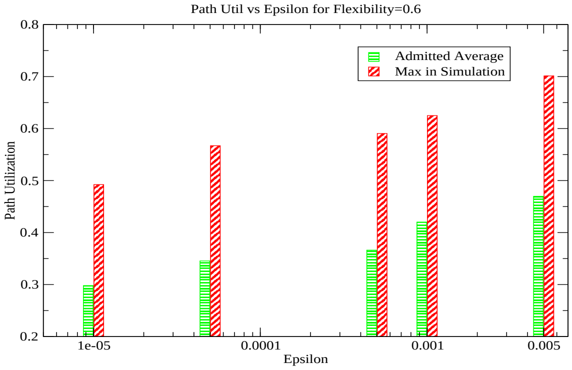

Fig. 15. If probability of capacity violation ($\epsilon$) is to be maintained at the same level for higher flexibility, number of admitted contracts decreases and hence the average utilization decreases (compare with Fig. 14).

utilization. Flexibility introduces an additional dimension to this aspect by allowing increase in the *maximum* achievable utilization for a given average utilization.

To illustrate this ability of flexibility, we turn to Fig. 13. For a fixed number of admitted contracts, the maximum path utilization increases with flexibility. It is important to note that in this case the number of admitted contracts has been held constant and not the violation probability. Consequently, higher flexibility allows higher utilization at the cost of a worse violation probability.

The role of $\epsilon$ in increasing the achievable average utilization is illustrated in Fig. 14 and Fig. 15. In this case the average utilization for a given $\epsilon$ as seen in Fig. 15 is lower as compared to Fig. 14 for a higher flexibility. This is because, the number of contracts admitted has to reduce to accommodate the same violation probability.

### *F. Effect of Bias in Traffic*

The admission control criterion rejected a contract if it violated the capacity constraint of even one path. If there is more demand toward certain destinations, i.e. the load is biased, there would be some resource wastage. Here we just present this effect and do not provide a solution.

We recall that in the simulations the quantities $m_{ij}$ were computed as $m_{ij} = r_i m_{1j}, i > 1$ where $r_1 = 1$, and $r_i, i > 1$ is a uniform random variable over $[max, min]$. Further $m_{1j}(\sum_i r_i) = 1$. If we increase $max$ we increase the bias of traffic toward certain destinations. Thus we obtain the number of admitted contracts and utilization for lower and higher bias cases in Fig. 16 and Fig. 17.

We see that the number of admitted contracts is lower for the same $\epsilon$ if the bias is higher. Similarly, the maximum measured utilization is lower in the case of higher bias in most cases.

## VI. SUMMARY AND CONCLUSIONS

This paper proposed a novel QoS architecture called the point-to-set architecture. The traditional point-to-point model was extended to be able to provide considerable freedom to the user network in dynamically apportioning the allocated bandwidth among a *finite set* of destinations. The model captured the statistical characteristics of per-destination load by first defining the per-destination traffic *fraction* as a random variable

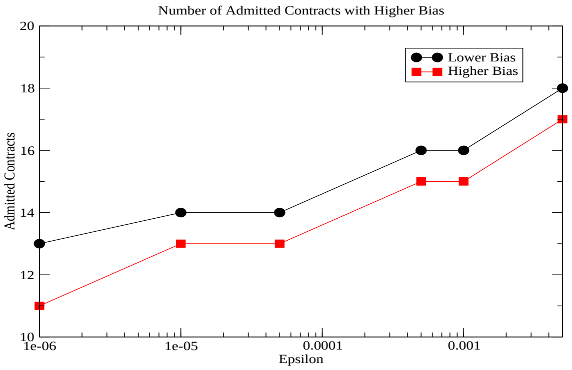

Fig. 16. Number of admitted contracts decreases with increase in *bias*. A higher bias indicates that a higher fraction of traffic is directed at a smaller subset of destinations

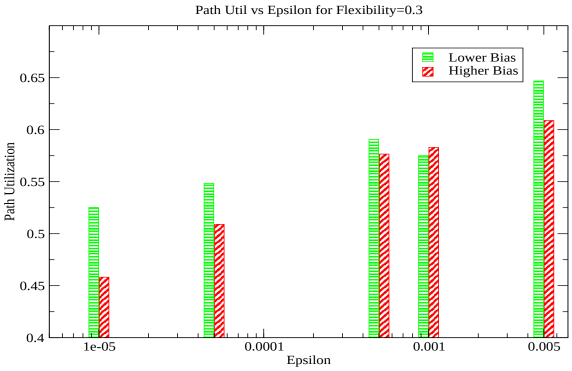

Fig. 17. Maximum measured utilization decreases with increase in *bias*. A higher bias indicates that a higher fraction of traffic is directed at a smaller subset of destinations

and second, letting the user specify a mean and variance for this random variable. The maximum permissible value for the ratio of this variance to the mean was defined to be the *flexibility* of the model. Exploiting the independence of user aggregates at the network edge nodes, a simple probabilistic admission control test was derived. The admission control procedure introduced the notion of a virtual path connecting the ingress to each destination egress with means to compute this path's capacity. The admission control test then involved computing the probability of violating any of the virtual path capacities.

The architecture was implemented in the NS-2 simulation environment and tested with real traffic traces. The simulation results demonstrated the superiority of the model over point-to-point models. The significance of flexibility and the permissible capacity violation probability ($\epsilon$) was characterized and verified by simulations. The parameters were shown to provide a control over the trade-off between multiplexing gains and loss rates.

Future work will involve studying means to further improve multiplexing gains by possible improvements to the admission control test. Reducing resource wastage when there is higher bias in offered load toward certain destinations also needs to be investigated.

## REFERENCES

[1] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang and W. Weiss, *An Architecture for Differentiated Services*, Dec. 1998, IETF RFC 2745.

[2] J.-Y. Le Boudec and P. Thiran, *Network Calculus*, Springer-Verlag, 2001.

[3] R. Braden, D. Clark and S. Shenker, *Integrated Services in the Internet Architecture: An Overview*, Jul. 1994, IETF RFC 1633.

[4] R.R. Boorstyn, A. Burchard, J. Liebeherr, C. Oottamakorn, "Statistical service assurances for traffic scheduling algorithms," *IEEE J. Select. Areas Commun.*, Vol. 18 No. 12, pp. 2651 -2664, Dec. 2000.

[5] C.S. Chang, "Stability, Queue Length, and Delay of Deterministic Stochastic Queuing Networks," *IEEE Trans. Automat. Control*, Vol. 39 No. 5, pp. 913-931, May 1994.

[6] D. Clark, W. Fang, "Explicit Allocation of Best-Effort Packet Delivery Service", *IEEE/ACM Trans. on Networking*, Vol. 6 No. 4, pp. 362-373, August 1998.

[7] B.T. Doshi, "Deterministic Rule Based Traffic Descriptors for Broadband ISDN: Worst Case Behavior and its Impact on Connection Acceptance Controls," *International Journal of Communication Systems*, March-April 1995

[8] N.G. Duffield, P. Goyal, A. Greenberg, P. Mishra, K.K. Ramakrishnan, J.E. van der Merive, "A Flexible Model for Resource Management in Virtual Private Networks", *ACM SIGCOMM 99*

[9] A. Elwalid, D. Mitra, R.H. 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, August 1995

[10] D. Ferrari, "Design and Application of a Delay Jitter Control Scheme for Packet-switching Internetworks," *Computer Communications*, 15(6):367-373, July 1992.

[11] G. Kesidis, T. Konstantopoulos, "Extremal Shape-Controlled Traffic Patterns in High-Speed Networks," *Trans. on Communications*, Vol. 48, No. 5, pp. 813-819, May 2000.

[12] E.W. Knightly, "Second Moment Resource Allocation in Multi-Service Networks," *ACM SIGMETRICS 97*, Vol 25 No. 1, pp. 181-191, Jun. 1997.

[13] E.W. Knightly, "Enforceable Quality of Service Guarantees for Bursty Traffic Streams," *IEEE INFOCOM 98*, Vol. 2, pp. 635-642, 1998.

[14] E.W. Knightly, N.B. Shroff, "Admission control for statistical QoS: theory and practice," *IEEE Network*, Vol. 13 No. 2, pp. 20-29, March 1999.

[15] A. Kumar, R. Rastogi, A. Silberschatz, B. Yener, "Algorithms for Provisioning Virtual Private Networks in the Hose Model", *ACM SIGCOMM 2001*, pp. 135 - 146, 2001.

[16] J. Kurose, "On Computing Per-session Performance Bounds in High-Speed Multi-hop Computer Networks," *Proc. of ACM SIGMETRICS 1992*, Vol. 20 No. 1, pp. 128-139, 1992.

[17] J. Liebeherr, S. Patek and E. Yilmaz, "Tradeoffs in designing networks with end-to-end statistical QoS guarantees," *Eighth International Workshop on Quality of Service*, pp. 221-230, 2000.

[18] F. Lo Presti, Zhi-Li Zhang, J. Kurose, D. Towsley, "Source Time Scale and Optimal Buffer/Bandwidth Tradeoff for Heterogeneous Regulated Traffic in a Network Node," *IEEE/ACM Trans. on Networking*, Vol. 7, No. 4, pp. 490-501, August 1999.

[19] J. Micheel, I. Graham, N. Brownlee, “The Auckland data set: an access link observed,” *Proceedings of the 14th ITC Specialists Seminar on Access Networks and Systems*, Catalonia, Spain, April 2001.

[20] NLANR “Auckland-IV trace archive,” http://pma.nlanr.net/Traces/long/auck4.html

[21] UCB/LBLN/VINT “Network Simulator - ns (version 2)”, http://www.isi.edu/nsnam/ns, 1997.

[22] S. Raghunath, K. Chandrayana, S. Kalyanaraman, “Edge-based QoS Provisioning for Point-to-Set Assured Services,” *IEEE ICC 2002*, Vol. 2, pp. 1128-1134, 2002.

[23] S. Rajagopal, M. Reisslein, K.W. Ross, “Packet Multiplexers with Adversarial Regulated Traffic,” *IEEE INFOCOM 98*, Vol. 1, pp. 347-355, March 1998.

[24] M. Reisslein, K.W. Ross and S. Rajagopal, “A Framework for Guaranteeing Statistical QoS,” *IEEE/ACM Trans. on Networking*, Vol 10, Issue 1, pp. 27-42, February 2002.

[25] V. Sivaraman and F. Chiussi, “Providing end-to-end statistical delay guarantees with earliest deadline first scheduling and per-hop traffic shaping,” *IEEE INFOCOM 2000*, Vol. 2, pp. 631-640, 2000.

[26] “Traffic Trace for ns”, http://www.research.att.com/˜breslau/vint/trace.html

[27] I. Stoica, H. Zhang, “LIRA: An Approach for Service Differentiation in the Internet,” *NOSSDAV 98*, Cambridge, England, pp. 115-128, July 1998.

[28] M. Vojnović and J.-Y. Le Boudec, “Stochastic Analysis of Some Expedited Forwarding Networks,” *IEEE INFOCOM 2002*, Vol. 2, pp. 1004-1013, 2002.

[29] H. Zhang and D. Ferrari, “Rate-Controlled Service Disciplines,” *Journal of High Speed Networks: Special Issue on Quality of Service*, 3(4), 1994.