Probabilistic Delay Guarantees Using Delay Distribution Measurement [*]

Kartik Gopalan
Florida State University
kartik@cs.fsu.edu

Tzi-cker Chiueh
Stony Brook University
chiueh@cs.sunysb.edu

Yow-Jian Lin
Telcordia Research
yjlin@research.telcordia.com

ABSTRACT

Carriers increasingly differentiate their wide-area connectivity offerings by means of customized services, such as virtual private networks (VPN) with Quality of Service (QoS) guarantees, or QVPNs. The key challenge faced by carriers is to maximize the number of QVPNs admitted by exploiting the statistical multiplexing nature of input traffic. While existing measurement-based admission control algorithms utilize statistical multiplexing along the bandwidth dimension, they do not satisfactorily exploit statistical multiplexing along the *delay dimension* to guarantee *distinct per-QVPN delay bounds*. This paper presents Delay Distribution Measurement (DDM) based admission control algorithm, the first measurement-based approach that effectively exploits statistical multiplexing along the delay dimension. In other words, DDM exploits the well known fact that the actual delay experienced by most packets of a QVPN is usually far smaller than its worst-case delay bound requirement since multiple QVPNs rarely send traffic bursts at the same time. Additionally, DDM supports QVPNs with distinct probabilistic delay guarantees – QVPNs that can tolerate more delay violations can reserve fewer resource than those that tolerate less, even though they require the same delay bound. A comprehensive performance evaluation using Voice over IP traces shows that, when compared to deterministic admission control, DDM can potentially increase the number of admitted QVPNs (and link utilization) by up to a factor of 3.0 even when the delay violation probability is as small as 10^{-5}.

Categories and Subject Descriptors: C.2.1 [Computer Communication Networks]: Network Architecture and Design

General Terms: Algorithms Measurement Performance

Keywords: Admission Control, Measurement-based, Statistical Multiplexing

1. INTRODUCTION

Performance-centric network applications such as Voice over IP (VoIP), video conferencing, online trading and streaming media, require dedicated network resources to meet their stringent delay and throughput requirements. An emerging service offering that meets this need is a Virtual Private Network (VPN) with QoS guarantees (or QVPN) that acts as a traffic trunk carrying aggregated traffic. Technologies such as Multi-Protocol Label Switched (MPLS) networks can map each QVPN to a long-term Label Switched Path (LSP). For instance, a QVPN could be a long-term Voice over IP (VoIP) trunk that carries aggregate traffic from several voice sessions rather than just one individual voice session. Thus QVPNs are set up and torn down over longer timescales and carry aggregate traffic that is less dynamic in nature than short-lived individual connections.

The key challenge faced by every network infrastructure provider is *to maximize the utilization efficiency of the network infrastructure and still support the stringent QoS requirements of each QVPN.* Maximizing utilization efficiency calls for an effective admission control algorithm that admits as many QVPNs as possible while allocating the least amount of resources needed to satisfy their QoS requirements. A simple approach of deterministic admission control allocates all the resources needed to ensure that the QoS guarantees are never violated. In the context of delay guarantees, this boils down to ensuring that packet delays never exceed the worst case delay bounds for each QVPN. On the flip side, worst case delays are rarely encountered in practice and a large proportion of network resources remain underutilized. The long-term and stable nature of each QVPN's aggregate real-time traffic provides the opportunity to improve network resource utilization by exploiting two specific statistical effects:

(i) Tolerance to delay violations: Most real-world real-time applications can tolerate a small fraction of excess delays or packet losses in their network traffic [25]. For instance, VoIP sessions can tolerate up to 10^{-3} fraction of their packets experiencing excess delays or losses without perceptually affecting audio quality. If 99.9% of the packets are observed to experience at most 50% of their expected worst-case delay, an admission control algorithm can potentially reserve only half of the resources that deterministic admission control would have reserved.

(ii) Statistical multiplexing along delay dimension: Due to statistical multiplexing, not all the QVPNs on a link experience their peak traffic bursts simultaneously. The con-

[*]This is an extended version of the original paper in ACM Multimedia 2004.

MM'04, October 10-16, 2004, New York, New York, USA.

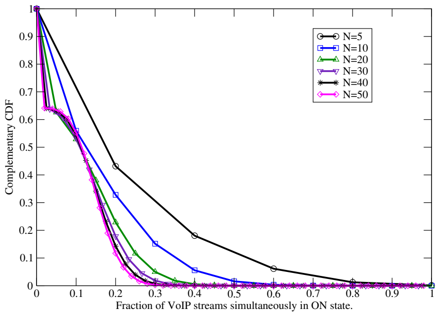

Figure 1: Complementary CDF of the fraction of VoIP sessions in ON state simultaneously as the number of VoIP sessions (N) in aggregate QVPN is varied.

sequence of this multiplexing is that packet delays rarely approach worst-case delays bounds that are based on all QVPNs transmitting at their peak burst simultaneously. To illustrate this multiplexing effect, we aggregated the ON–OFF packet traces for different number of recorded VoIP sessions (described later in Section 6). Figure 1 shows the complementary cumulative distribution function (CCDF) of the fraction of VoIP sessions in an aggregate that are simultaneously in their ON state. We observe that half the time less than 12% of the VoIP sessions are in their ON state simultaneously and its almost never the case that more than 40% of the sessions are simultaneously active. Similar statistical effects can be expected for other categories of real-time network traffic such as video conferencing and online financial transactions.

This paper proposes a practical and efficient link-level measurement-based algorithm, called *Delay Distribution Measurement* (DDM) based admission control, that exploits the above two statistical effects to maximize the number of QVPNs admitted with performance guarantees. The QoS parameters that the DDM algorithm supports include delay bound, delay violation probability bound and the long-term average bandwidth. DDM is the first measurement-based algorithm that simultaneously provides all the following features.

- **Statistical multiplexing along delay dimension:** DDM is the first measurement-based approach which exploits statistical multiplexing along the delay dimension to increase resource utilization in comparison to purely deterministic admission control. In contrast, the earlier measurement-based approaches mainly focused on statistical multiplexing along the bandwidth dimension, i.e. multiplexing due to the fact that QVPNs often transmit at rates much below their stated long-term bandwidth requirement.

- **Distinct per-QVPN probabilistic delay bounds:** DDM supports QVPNs for which a certain percentage of delay bound violations are tolerable. The key difference from prior approaches is DDM's ability to differentiate among QVPNs in terms of their tolerance to delay bound violations. QVPNs with higher tolerance to delay bound violations are allocated fewer resources than those with lower tolerance, even though they may have the same delay bound requirement.

- **Unified support for probabilistic and deterministic delay bounds:** DDM provides a single admission control framework to support QVPNs that may have probabilistic or deterministic delay bounds. Deterministic delay bound requirements simply correspond to zero tolerance to delay violations.

The principal challenge in providing per-QVPN probabilistic delay guarantees is to determine the mapping between delay bound, delay violation probability bound and resource requirements. DDM dynamically measures the service delay of each packet, computes the ratio between the actual service delay and the worst-case delay that the packet could experience, and derives a delay ratio distribution. This dynamically measured delay ratio distribution is used to derive the bandwidth reservation needed to support a given probabilistic delay bound. Once the DDM algorithm reserves an amount of bandwidth for a QVPN, a rate-based packet-by-packet scheduler (such as WFQ [19] or Virtual Clock [26]) guarantees the assigned bandwidth share.

2. RELATED WORK

The principal features that distinguish DDM from earlier works are (1) its ability to exploit statistical multiplexing along the delay dimension, in contrast to bandwith dimension, and (2) its ability to provide a distinct probabilistic delay guarantee to each QVPN, as opposed to shared guarantees in earlier approaches. The literature on exploiting statistical multiplexing is extensive and we discuss the ones most relevant to this work. Knightly and Shroff [16] provide an excellent overview of admission control approaches for link-level statistical QoS.

Kurose [17] derived probabilistic bounds on delay and buffer occupancy of QVPNs using the concept of stochastic ordering for network nodes that use FIFO scheduling. Unlike FIFO schedulers that inherently cannot differentiate between performance requirements of different QVPNs, we are interested in real-time traffic schedulers that can provide per-QVPN delay and bandwidth guarantees. Reisslein et al. [21] have derived statistical delay bounds for traffic in a single link and network settings using a fluid traffic model. Their work approximates the loss probability at a link using independent Bernoulli random variables. All QVPNs share a common buffer with traffic loss assumed to be split among QVPNs in proportion to their input rates. In contrast, we assume a packet-based model, an independent buffer space for each QVPN, and permit explicit specification of delay violation probability bound for each QVPN. Elwalid and Mitra[7] have proposed a scheme to provide statistical QoS guarantees in the GPS service discipline for two guaranteed traffic classes and one best effort class. Again a fluid traffic model was considered. Le Boudec and Vojnovic [18] consider stochastic delay guarantees in expedited forwarding (EF) networks with aggregate scheduling. Their work operates under the Diffserv framework in which EF traffic is marked at the network ingress. Each forwarding node in the network interior is abstracted by a service curve and provides a common stochastic rate and latency guarantee to all transiting EF traffic. Several analytical approaches [15, 4, 6, 11] have also considered the performance of multiplexing with a shared buffer for independent regulated inputs. In contrast, we consider distinct per-QVPN probabilistic delay bounds with independent buffers for independent regulated inputs. Schemes for providing probabilistic QoS in networks using Earliest Deadline First (EDF) scheduling were proposed in [1, 22, 2]. Unlike the rate-based schedulers considered here, EDF decouples rate and delay guarantees at the expense of admission control complexity. Additionally, it is difficult to guarantee distinct per-QVPN delay violation probabilities with EDF due to strong interactions among QVPNs sharing a link. In contrast, rate-based schedulers, such as the one we use, provide explicit performance isolation among QVPNs and are especially suited to guarantee QVPN-specific delay violation probabilities.

Several existing measurement based admission control algorithms (MBAC) address QoS requirements along the dimensions of the bandwidth or aggregate loss-rate. The notion of Effective Bandwidth [14] is an important concept in MBAC algorithms that provides a measure of bandwidth resource usage by flows relative to their peak and mean usage. Breslau *et. al.* [3] performed a comparative study of several MBAC algorithms [20, 12, 8, 9, 5] under FIFO service discipline and concluded that none of them are capable of accurately achieving loss targets. Qiu and Knightly[20] proposed an MBAC scheme that measures maximal rate envelopes of aggregate traffic. An important difference of our algorithm with the existing MBAC schemes is that the latter mainly focus on providing bandwidth guarantees but do not address QVPNs that require distinct statistical delay bounds on a per-QVPN basis.

In a different context of shared application hosting platforms, a measurement-based approach is presented in [23] to support heterogeneous applications with probabilistic bandwidth requirements. In a sense, the technique used for measurement in [23] can be considered as being closest to ours. Their approach performs offine profiling of each individual application's bandwidth usage distribution to derive probabilistic bandwidth requirements. DDM relies on aggregate delay distribution curve as well for deriving resource reservations and admission control. However, our delay distribution curve represents aggregate run-time delay distribution of all QVPNs, in contrast to individual offline measurement of bandwidth distribution for each application.

3. WORST-CASE DELAY BOUND

In this section, we review some classical results for delay bounds using rate-based schedulers. A QVPN F_i is defined as an aggregate that carries traffic with an average bandwidth of ρ_i^{avg} and burst size σ_i. We consider the context of a single link l with capacity C_l. Traffic belonging to every QVPN F_i that traverses link l requires each packet to be serviced by the link scheduler within a delay bound $D_{i,l}$ and with a delay violation probability no greater than $P_{i,l}$. For instance, if $D_{i,l} = 10ms$ and $P_{i,l} = 10^{-5}$, it means that no more than a fraction 10^{-5} of packets belonging to the QVPN can experience a delay greater than $10ms$. We assume that each QVPN's incoming traffic is regulated by a token bucket with bucket depth σ_i and token rate ρ_i^{avg}. The amount of QVPN F_i traffic arriving at the scheduler in any time interval of length τ is bounded by $(\sigma_i + \rho_i^{avg}\tau)$. [1]

The job of a link scheduler is to prioritize the transmission of packets belonging to different QVPNs over a common link. We assume that packets are serviced by *rate-based link schedulers*, such as WFQ [19] or Virtual Clock [26]. It can be shown that the worst-case queuing delay $D_{i,l}^{wc}$ experienced at a link l by any packet belonging to a QVPN F_i under the WFQ service discipline is given by the following expression.

$$D_{i,l}^{wc} = \frac{\sigma_i}{\rho_{i,l}} + \frac{L_{max}}{\rho_{i,l}} + \frac{L_{max}}{C_l} \quad (1)$$

where σ_i is F_i's burst size at link l, L_{max} is the maximum packet size, $\rho_{i,l}$ is the reservation for F_i at link l, and C_l is the total capacity of link l. The first component of the delay is fluid fair queuing delay, the second component is the packetization delay and the third component is scheduler's non-preemption delay. We are interested in rate-based schedulers since, in their case, the relationship between delay bound and the amount of bandwidth reserved for a QVPN can be explicitly specified. Furthermore, as we will see in Section 4, rate-based schedulers enable us to differentiate among QVPNs in terms of their delay violation probability requirements. In contrast, for non rate-based schedulers, such as Earliest Deadline First (EDF), the resource-delay relationship is difficult to determine, which in turn makes the admission control process more complicated. Hence, even though non rate-based schedulers can potentially provide higher link utilization, it is difficult to guarantee delay violation probability on a *per-QVPN basis*.

4. DELAY TO RESOURCE MAPPING

Probabilistic delay guarantees assist in reducing the bandwidth reservation for each QVPN by exploiting their tolerance to certain level of delay violations. The statistical multiplexing effect ensures that bursts of size σ_i from different QVPNs F_i tend to be temporally spread out and rarely occur at the same time. As a result, worst-case delay is rarely experienced by packets traversing a link. Assume that the request for a QVPN F_i specifies its average rate ρ_i^{avg}, burst size σ_i, required delay bound $D_{i,l}$ and delay violation probability $P_{i,l}$ at link l. Each QVPN F_i traversing the link is assigned a bandwidth reservation $\rho_{i,l} \geq \rho_i^{avg}$, which satisfies both the delay requirement $(D_{i,l}, P_{i,l})$ as well as the average rate requirement ρ_i^{avg}. Note that ρ_i^{avg} is the long-term average rate of F_i whereas the bandwidth reservation $\rho_{i,l}$ is used by the scheduler to determine run-time preference for F_i's traffic over other QVPNs. In this section, we derive the correlation function that maps F_i's specification $(\rho_i^{avg}, \sigma_i, D_{i,l}, P_{i,l})$ to its bandwidth reservation $\rho_{i,l}$.

4.1 CDF Construction

Assume that for each packet k, the system tracks the run-time measurement history of the ratio r_k of the actual packet delay experienced $D_{i,l}^k$ to the worst-case delay $D_{i,l}^{wc}$, i.e., $r_k = D_{i,l}^k / D_{i,l}^{wc}$ where r_k ranges between 0 and 1. We can use these measured samples of ratio r_k to construct a cumulative distribution function (CDF) $Prob(r)$. The distribution $Prob(r)$ gives the probability that the ratio between the actual delay encountered by a packet and its worst-case delay is smaller than or equal to r. Conversely, $Prob^{-1}(p)$ gives the maximum ratio of actual delay to worst-case delay that can be guaranteed with a probability p. Figure 2 shows an example of a CDF constructed in this manner for a

[1] Our algorithm remains essentially unchanged in case of dual/cascaded leaky bucket regulators as well.

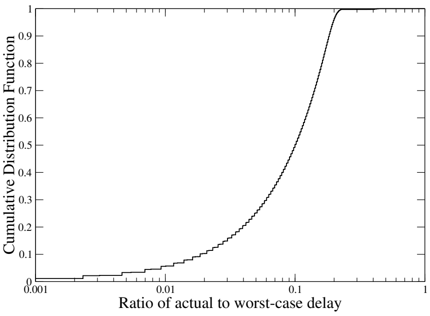

Figure 2: Example of cumulative distribution function (CDF) of the ratio of actual delay to worst-case delay experienced by packets. X-axis is in log scale to highlight the ratio distribution in the low-ratio range. 39 VoIP QVPNs traverse a 10Mbps link. $\rho_i^{avg} = 256$Kbps. Delay bound=10ms. Delay violation probability=10^{-5}.

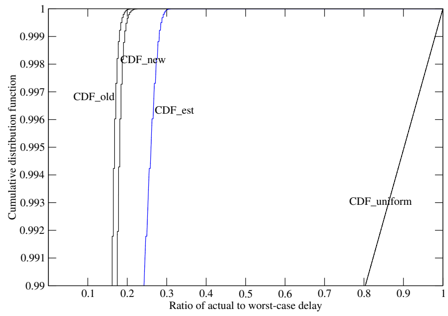

Figure 3: Example of different CDF curves for one simulation scenario. X-axis is in linear scale to highlight the difference between measured and estimated CDF curves. The Y-axis range shown is from 0.99 to 1.0 which corresponds to the typical tolerance range for delay violations (below 10^{-2}).

specific simulation scenario of 39 VoIP QVPNs. (Simulation details follow in Section 6).

To construct the CDF in practice, we partition the ratio range from 0 to 1 into a number of sub-ranges, and then for each sub-range, keep updating the count of packets transmitted whose ratio r_k falls within the sub-range. The CDF can be constructed by computing the accumulated count of packets from the lowest sub-range to each sub-range i.

The CDF would typically be maintained over a sliding measurement window.

4.2 Resource Mapping

The CDF curve $Prob(r)$ concisely quantifies the level of statistical multiplexing along the delay dimension. For instance, Figure 2 indicates that most of the packets experience less than 1/4th of their expected worst-case delay. Thus, reserving resources to cover for the worst-case delay is wasteful since it is rarely encountered in practice. In this section, we describe how we can exploit the statistical multiplexing information quantified by $Prob(r)$, in addition to each QVPN's tolerance to delay violations, to reduce the amount of per-QVPN bandwidth reservation.

Given the measured estimate of functions $Prob(r)$ and $Prob^{-1}(p)$, the following expression determines the delay-derived bandwidth reservation $\rho_{i,l}^{delay}$ required to satisfy QVPN F_i's probabilistic delay requirement $(D_{i,l}, P_{i,l})$.

$$D_{i,l} = \left(\frac{\sigma_i + L_{max}}{\rho_{i,l}^{delay}} + \frac{L_{max}}{C_l} \right) \times Prob^{-1}(1 - P_{i,l}) \quad (2)$$

Equation 2 states that in order to obtain a delay bound of $D_{i,l}$ with a delay violation probability bound of $P_{i,l}$, we need to reserve a minimum bandwidth of $\rho_{i,l}^{delay}$ which can guarantee a worst-case delay of $D_{i,l}^{wc} = D_{i,l}/Prob^{-1}(1-P_{i,l})$. Conversely, the delay-derived bandwidth requirement $\rho_{i,l}^{delay}$ of a QVPN F_i at link l is

$$\rho_{i,l}^{delay} = \mathcal{B}_l(D_{i,l}, P_{i,l}, \sigma_i) = \frac{\sigma_i + L_{max}}{\frac{D_{i,l}}{Prob^{-1}(1-P_{i,l})} - \frac{L_{max}}{C_l}} \quad (3)$$

The actual reservation required to satisfy QVPN F_i's QoS requirement $(\rho_{i,l}^{avg}, \sigma_i, D_{i,l}, P_{i,l})$ is $\rho_{i,l} = max\{\rho_{i,l}^{avg}, \rho_{i,l}^{delay}\}$. In other words, the actual bandwidth reservation for a QVPN is dictated by the tighter of two QoS requirements - one imposed by its average bandwidth requirement $\rho_{i,l}^{avg}$ and the other imposed by its probabilistic delay requirement $(D_{i,l}, P_{i,l})$. It is worth pointing out once more that the above resource mapping function exploits statistical multiplexing along the delay dimension, rather than along the bandwidth dimension as in earlier approaches. This is a direct consequence of the fact that we measure the distribution of actual to worst-case delay ratio (in contrast to distribution of bandwidth usage as in [23]). Specifically, if $\rho_{i,l}^{delay}$ happens to be larger than $\rho_{i,l}^{avg}$ for all QVPNs then the resource allocation will be guided by statistical delay requirements rather than deterministic bandwidth requirements.

5. ADMISSION CONTROL USING DDM

In this section, we describe the DDM admission control algorithm for admitting a new QVPN F_N that arrives at a link l on which $N-1$ QVPNs have already been admitted. The principal challenge of admission control lies in estimating the impact of F_N's traffic on guarantees provided to already admitted QVPNs. If the F_N is admitted, it will result in an increase in traffic load carried by the link and consequently larger actual delays experienced by packets from all QVPNs. Specifically, the CDF of actual to worst-case delay ratio will tend to become more conservative by shifting to the right after F_N becomes active. Hence it is important that, even before F_N can be admitted, DDM must estimate and account for the impact of the new QVPN on the delay distribution of existing QVPNs.

The DDM algorithm consists of two phases. The first phase estimates the expected delay distribution assuming QVPN F_N is admitted. The second phase performs the actual admission control using the estimated CDF from first phase and computes future resource requirements of all QVPNs (including the new one). F_N is admitted only if each QVPN's resource requirement can be satisfied within the available link capacity.

5.1 Significance of CDF Evolution

If the new QVPN F_N is admitted, the link with a finite capacity C_l has to shoulder the additional traffic load from F_N. As a result packets for all QVPNs traversing the link will experience larger delays on the average. More specifically, the additional load from F_N could impact the CDF curve shown in Figure 2 by shifting it to the right. In other words, for the same delay violation probability p, if $r_1 = Prob_{old}^{-1}(1-p)$ before admitting F_N and $r_2 = Prob_{new}^{-1}(1-p)$ after admitting F_N, then $r_2 \geq r_1$. Because a larger value of $Prob_{new}^{-1}(1-p)$ translates into higher bandwidth requirement in Equation 3, CDF_{new} is said to be more conservative than CDF_{old} since CDF_{new} can admit fewer QVPNs than CDF_{old}. Figure 3 provides an example of CDF_{old} and right-shifted CDF_{new} for one simulation scenario in the Y-axis range from 0.99 to 1.0 (since this range happens to be of most interest).

If we simply use CDF_{old} to derive the bandwidth reservation for F_N, and the actual CDF_{new} turns out to be significantly more conservative than CDF_{old}, F_N may be assigned a much smaller bandwidth than what it actually needs to meet its probabilistic delay requirement. The key research challenge of DDM algorithm thus lies in how to predict the impact of the new QVPN F_N on the delay distribution of $(N-1)$ existing QVPNs without assuming any apriori traffic model.

The impact of new QVPN F_N on CDF_{old} depends upon several factors. In general tight QoS requirements - such as a small delay requirement $D_{N,l}$, a low tolerance to delay violation $P_{N,l}$, a large average rate ρ_N^{avg} or a big burst size σ_N – all lead to larger ratio of actual to worst-case delay and a more conservative CDF. Furthermore, the increment from $Prob_{old}^{-1}(1-p)$ to $Prob_{new}^{-1}(1-p)$ could be different for different values of violation probability p. Finally, the magnitude of a new QVPN's relative load contribution to a link's traffic affects the amount of difference between the CDFs before and after the new QVPN is admitted.

5.2 Predicting CDF Evolution

Given the multitude of factors that influence the evolution of CDF, it is difficult if not impossible, to exactly predict CDF_{new} using CDF_{old} and QVPN F_N's QoS requirements. The DDM algorithm uses a heuristic approach to approximate CDF_{new}. Let τ be the length of a moving time window over which the delay distribution CDF_{old} of existing $N-1$ QVPNs is measured. Let m be the number of packets generated by $N-1$ QVPNs that traverse the link in duration τ. In a time interval τ, F_N can potentially transmit a maximum of $n = \sigma_N/L_{min} + \rho_N^{avg} * \tau/L_{min}$ number of packets, where L_{min} is the minimum packet size. Assume that these n additional packets experience a uniform distribution of actual to worst-case delay ratio. A uniform distribution is a very conservative estimate of delay distribution (though not the most conservative one) which assumes that packet delays for the new QVPN F_N are expected to be uniformly distributed over the range of ratios from 0 to 1 and that all packets are of size L_{min}. In reality, a large majority of packets experience small packet delays (as seen in Figure 2) and are of size greater than L_{min}.

To characterize CDF_{new}, we first combine the uniform delay ratio distribution for F_N obtained above with a weight of $\frac{n}{n+m}$ and the delay ratio distribution CDF_{old} with a weight of $\frac{m}{n+m}$ to obtain a distribution called $CDF_{uniform}$, which represents an estimate of the cumulative distribution that would result if F_N were fully loaded and the delay ratio of the packets from F_N were distributed uniformly between 0 and 1. $CDF_{uniform}$ can be constructed using the technique described in Section 4, but with the difference that before computing the accumulated sum for each ratio sub-range, we add n/R to the count of ratio samples in each sub-range, where R is the number of sub-ranges between 0 and 1. In other words, n delay ratios are assumed to be uniformly distributed over all ratio sub-ranges.

Empirically $CDF_{uniform}$ is a very conservative estimate of the distribution CDF_{new}, because both the uniform delay ratio distribution assumption and the full load assumption are too pessimistic. As a result, CDF_{new} lies somewhere between CDF_{old} and $CDF_{uniform}$ constructed above. We further approximate CDF_{new} by constructing CDF_{est}, which in turn is a weighted combination of CDF_{old} and $CDF_{uniform}$. Specifically,

$$Prob_{est}^{-1}(1-p) = \alpha Prob_{uniform}^{-1}(1-p) + (1-\alpha) Prob_{old}^{-1}(1-p) \quad (4)$$

The factor α is the *impact factor* that determines how far the distribution curve CDF_{est} is from $CDF_{uniform}$ and CDF_{old}. For a new QVPN that imposes a relatively large load on the link with respect to existing load, CDF_{est} should be close to $CDF_{uniform}$ since the latter is more conservative in admitting QVPNs. On the other hand, for a new QVPN that imposes a relatively small load with respect to existing load, CDF_{est} should be closer to CDF_{old} since in this case the new QVPN has a relatively smaller impact on CDF_{old}. With this consideration in mind, we define the impact factor as the fraction of new QVPN F_N's load on the total expected load.

$$\alpha = \frac{\rho_{N,l}}{\sum_{i=1}^{N} \rho_{i,l}} \quad (5)$$

Here $\rho_{i,l}$ is computed using the distribution $CDF_{uniform}$ since it is the only estimate of future delay distribution we have at the time of admitting F_N. Since we are practically interested in only the delay violation probabilities $P_{i,l}$ for existing and new QVPNs, we only need to compute that portion of CDF_{est} which covers these delay violation probabilities of our interest; typically the violation probabilities lie in the range 10^{-2} to 10^{-6} which corresponds to a small upper portion of the Y-axis in Figure 2. An example of different CDF curves is illustrated in Figure 3 within the Y-axis range of 0.99 to 1 for one simulation scenario. We see that CDF_{est} is the closest approximation to CDF_{new}, although a bit more conservative. $CDF_{uniform}$ is the most conservative of all.

Note that constructing CDF_{est} involves two levels of weighted combinations - first in constructing $CDF_{uniform}$ from CDF_{old} and a uniform distribution of new QVPN's packets, and second in constructing CDF_{est} from CDF_{old} and $CDF_{uniform}$. The difference is that the $CDF_{uniform}$ provides a first-cut conservative estimate of CDF_{new} whereas this estimate is further refined by constructing CDF_{est}. In Section 6, we validate that the above technique for CDF estimation indeed reliably captures the future delay distribution of admitted QVPNs.

5.3 The Admission Control Algorithm

With the delay-probability-bandwidth correlation function in place, we now present the DDM admission control algorithm in Figure 4. The algorithm can be invoked either to admit a new QVPN F_N or to periodically re-calculate the requirements of already admitted QVPNs. Without loss

Input : (a) $(D_{i,l}, P_{i,l}, \rho_i^{avg}, \sigma_i)$ for each QVPN F_i, $1 \le i \le N$.
(b) The measured delay ratio distributions.

Compute CDF_{old} and $CDF_{uniform}$ from delay ratio distributions.

for $i = 1$ to N
Compute $\rho_{i,l}^{delay} = \mathcal{B}_l(D_{i,l}, P_{i,l}, \sigma_i)$ using Equations 3 and 4.
$\rho_{i,l} = max\{\rho_i^{avg}, \rho_{i,l}^{delay}\}$

/*Perform admission checks*/
if $(\sum_{i=1}^{N} \rho_{i,l} > C_l)$ then
Reject QVPN F_N and exit.

/*QVPN F_N can be admitted*/
for $i = 1$ to N
Reserve bandwidth $\rho_{i,l}$ for F_i.

Figure 4: The DDM algorithm to determine whether a new QVPN F_N can be admitted such that each QVPN F_i, $1 \le i \le N$, can be guaranteed a delay bound $D_{i,l}$, delay violation probability $P_{i,l}$, and average rate ρ_i^{avg}.

of generality, the following discussion assumes the first scenario.

Assume that $N-1$ QVPNs are currently being served by the scheduler and F_N arrives for admission. The algorithm first calculates $CDF_{uniform}$ using the measured delay distribution CDF_{old} and QVPN F_N's average rate requirement ρ_N^{avg}. For each of the N QVPNs (including the new one) the algorithm next computes the delay-derived bandwidth requirement $\rho_{i,l}^{delay}$ using Equations 3 and 4. The actual bandwidth requirement $\rho_{i,l}$ is the larger of the delay-derived requirement $\rho_{i,l}^{delay}$ and average requirement $\rho_{i,l}^{avg}$. The new QVPN F_N is admitted only if following condition is satisfied.

$$\sum_{i=1}^{N} \rho_{i,l} \le C_l \qquad (6)$$

Equation 6 states that the sum of bandwidth requirements of all QVPNs, under the estimated delay ratio distribution CDF_{est}, should be smaller than C_l. The QVPN F_N is rejected if this condition cannot be satisfied. If the new QVPN is accepted, the algorithm sets the bandwidth reservation for each QVPN to $\rho_{i,l}$ as computed above.

The robustness of DDM algorithm, in essence, depends upon the accuracy of estimating CDF_{est} before admitting a new QVPN F_N. This is because the act of admitting F_N results in altering the reservation $\rho_{i,l}$ of already admitted flows F_1 to F_{N-1}. A CDF_{est} that is too conservative can lead to under utilization of a link's resources whereas one that is over optimistic can lead to a potential violation of QoS guarantees for all QVPNs at run time. The principal challenge in DDM algorithm lies in accurately estimating CDF_{est} before admitting F_N using an appropriate value of the impact factor α in Equation 4 – a value that is neither too optimistic nor too conservative. In our experience with experiments described in Section 6, an impact factor given in Equation 5, that equals the fractional load imposed by the new flow, provided a good estimate of CDF_{est}.

The admission control algorithm described above provides a unified framework to support QVPNs with both probabilistic as well as deterministic delay requirements. Specifically, QVPNs requiring deterministic delay bounds can simply be treated as requiring a violation probability of zero, which in turn can be easily factored into the calculation of $\rho_{i,l}$ described in Section 4.2. This is in part feasible due to the fact that DDM exploits statistical multiplexing along the delay dimension, but not along the bandwidth dimension. If statistical bandwidth multiplexing is to be supported as well, such as by means of over-subscription of link capacity, then QVPNs requiring deterministic guarantees may need explicit shielding from statistically multiplexed traffic.

5.4 Algorithm Complexity

The step for computing CDF_{old} and $CDF_{uniform}$ has $O(R)$ time complexity where R is the number of sub-ranges in the delay ratio interval from 0 to 1. The subsequent steps in the algorithm have $O(N)$ time complexity where N is the number of QVPNs being considered. Thus the complexity of the DDM algorithm is $O(N+R)$. In practice, the first step of computing CDF_{old} and $CDF_{uniform}$ is the more dominant of the two components due to the larger number of sub-ranges R. The algorithm itself is invoked quite infrequently only when either new QVPN requests arrive for admission at the link or existing QVPN reservations need to be periodically recomputed. The run-time computation overhead of maintaining CDFs is also minimal since we only need few arithmetic operations to record the ratio for each packet transmitted by the link.

In terms of space cost, the only significant additional space required is in the order of $O(R)$ (about 400KB with $R =$ 100K) for maintaining CDF_{old}, which represents aggregate delay distribution information for all QVPNs. The values for $CDF_{uniform}$ and CDF_{est} can be derived as and when required during admission control. In particular, DDM requires no additional space for per-QVPN state maintenance when compared to any other algorithm that provides per-QVPN QoS. In our context, QVPNs represent a limited number of traffic aggregates (such as LSPs in MPLS), rather than individual TCP/IP connections, which further reduces the space requirement to within practical bounds.

6. PERFORMANCE OF DDM

In this section, we study the performance of the DDM algorithm in comparison to deterministic admission control. We use the deterministic approach, instead of one of the earlier approaches, as a baseline due to the following reasons. First, earlier measurement-based approaches mainly address multiplexing along the bandwidth dimension, i.e. multiplexing due to the fact that QVPNs typically transmit at rates much below their stated long-term bandwidth requirement. In contrast, DDM exploits multiplexing along an orthogonal delay dimension which occurs even when individual QVPNs transmit at their stated bandwidth, i.e. multiplexing due to the fact different QVPNs transmit their traffic bursts at different times. Secondly, to the best of our knowledge, earlier analytical approaches that address probabilistic delay guarantees either assume a fluid traffic model (as opposed to a packetized model), or do not support distinct per-QVPN probabilistic delay bounds, but rather provide shared guarantees such as by multiplexing QVPN traffic in a shared buffer. Thus the problem addressed by DDM is fundamentally different from earlier approaches, and leaves deterministic admission control, rather than the earlier statistical or measurement-based approaches, as the baseline for comparison.

The real traffic traces used in our simulations are principally composed of VoIP sources due to the lack of representative traffic traces from other real-time applications such as video conferencing and online financial trading, However, a note regarding applicability of DDM to heterogeneous real-time traffic is in order. Unlike voice, video conferencing applications have relatively higher and more variable data rates (due to quantization via motion vectors and prediction algorithms), though with similar latency requirements. Online trading applications, on the other hand have much lower data rates with tighter latency requirements. In the presence of different categories of real-time traffic, we still expect significant potential gains in link utilization with varying degrees of statistical multiplexing. However, DDM algorithm is equally applicable to mixes of all categories of real-time traffic and nothing in the algorithm precludes any specific traffic category.

6.1 Evaluation Setup

Using the *ns-2* simulator, we configured a single link at 10 Mbps and packets arriving at the link were served by a WFQ scheduler. Each QVPN traffic consisted of aggregated traffic traces of recorded VoIP conversations used in [13], in which spurt-gap distributions were obtained using G.729 voice activity detector. Each VoIP stream had an average data rate of around 13 Kbps, peak data rate of 34 Kbps, and packet size of $L_{max} = 128$ bytes. We temporally interleaved the 20 VoIP streams to generate aggregate traffic trace for each QVPN with an aggregate data rate of $\rho_i^{avg} = 256$Kbps.

Each aggregated VoIP trace was 8073 seconds long. Every QVPN in our simulations sent traffic for the entire lifetime of the simulation with the aggregate traffic trace being repeated over its lifetime. Traffic from each admitted QVPN passed a token bucket with bucket depth of 1280 bytes (10 packets) and token rate of 256 Kbps. Each new QVPN required a guarantee on delay bound and a delay violation probability. The admission control algorithm decided whether to admit or reject the QVPN and how much bandwidth to reserve according to algorithm in Figure 4. Each QVPN was generated with a periodic inter-arrival time of $10,000$ seconds. The reason we selected periodic instead of exponential inter-arrival times (as in other works) is that our QVPNs are long lived and are expected to arrive fairly infrequently, so that the measured CDF can stabilize before being used to admit another QVPN. Hence the request arrival pattern does not significantly impact the admission control decisions. The CDF was measured over a time interval of $10,000$ seconds between QVPN arrivals. Each simulation run lasted for $1000,000$ seconds.

For simulations, we recorded the ratio of actual to worst-case delay of every packet traversing the link within the current CDF window (although in a realistic scenario an intelligent sampling mechanism would be more desirable). The observed ratios are accumulated into a histogram. The actual CDF is computed from the histogram only when making admission decisions or re-calculating existing reservations.

6.2 Per-QVPN Probability Bounds

We start by validating that the DDM algorithm can indeed provide distinct guarantees on heterogeneous delay violation probabilities for a mix of different traffic types. In the first experiment we consider a traffic mix in which all QVPNs request the same delay bound of 20ms, same average rate of 256Kbps, and same burst size of 10 packets, but require different guarantees on delay violation probability, the requirement being uniformly distributed among the four values 10^{-2}, 10^{-3}, 10^{-4} and 10^{-5}. Figure 5 plots the actual fraction of packets exceeding their delay bound against the desired violation probability for each QVPN that experiences any excess delay. The figure includes data points from 5 simulation runs with different random seeds and each data point represents the rate of delay violation experienced by one QVPN. Figure 6 plots the same data when the constituent QVPNs have heterogeneous delay bounds (10ms–30ms), data rates (256Kbps–2Mbps) and burst-sizes (10–40 packets), in addition to heterogeneous violation probability requirements (10^{-2}–10^{-5}). The line through the graph marks the limit above which the actual rate of delay violations would exceed the desired delay violation probability. The fact that all data points are below the line indicates that the actual delay violation rate is smaller than the maximum permissible for each QVPN. Furthermore the figure shows that QVPNs that have higher tolerance to delay violations are more likely to experience a higher rate of violation than QVPNs with lower tolerance. The DDM algorithm is able to distinguish among QVPNs in terms of delay violation rates

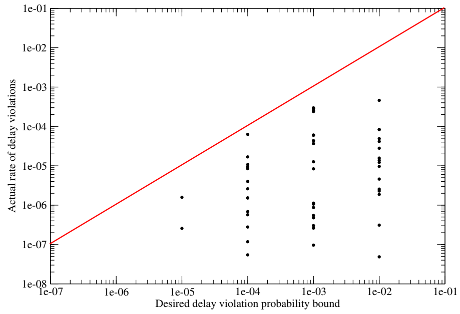

Figure 5: The DDM algorithm satisfies distinct per-QVPN delay violation guarantees, other requirements being the same. Each data point corresponds to one QVPN. Delay bound=20ms. Burst Size=10pkts. Link capacity = 10Mbps. The plot includes data points from 5 runs with different random seeds.

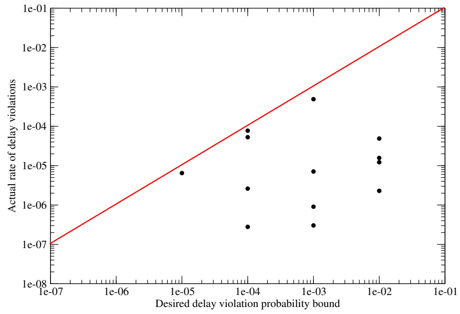

Figure 6: The DDM algorithm satisfies distinct per-QVPN delay violation guarantees even when constituent QVPNs have dissimilar delay bound, data rate, and burstiness requirements. Link capacity=10Mbps. Each data point corresponds to one QVPN. The plot includes data points from 5 runs with different random seeds.

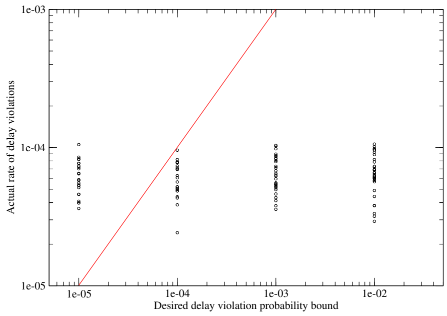

Figure 7: The pure over-subscription based algorithm cannot satisfy distinct per-QVPN delay violation guarantees. Each data point corresponds to one QVPN. Delay bound=20ms. Burst size=10pkts. Link capacity=10Mbps and is over-subscribed by a factor of 2.0. The plot includes data points from 5 simulation runs with different random seeds.

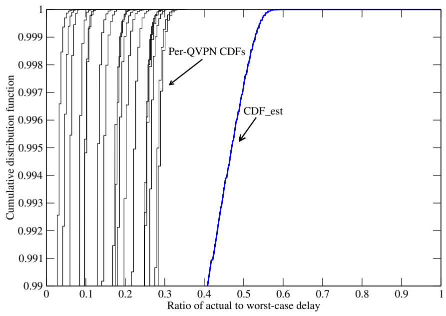

Figure 8: The predicted CDF_{est} (the right-most curve) indeed provides a reliable bound on future delay ratio distribution for each admitted QVPN (all other curves). The figures plots the ratio distribution from one representative simulation scenario. The Y-axis range shown (from 0.99 to 1.0) corresponds to the typical tolerance range for delay violations (below 10^{-2}).

because it assigns service bandwidth $\rho_{i,l}$ to QVPNs in the inverse proportion of their tolerance to delay violations. This translates to higher dynamic preference for packets belonging to QVPNs with low delay tolerance and vice-versa.

In the next experiment, we show that pure over-subscription of link capacity cannot provide distinct guarantees on heterogeneous delay violation probabilities. We use the same parameters as the previous experiment, except that instead of using the DDM algorithm, we use deterministic admission control and over-subscribe the link capacity by a factor of 2.0 so as to admit the same number of QVPNs as the DDM algorithm (i.e. 35 QVPNs) with no over-subscription. Figure 7 shows that irrespective of desired delay violation bounds, all QVPNs experience similar rates of actual delay violations. In fact, QVPNs with low tolerance (10^{-5}) to delay violations can experience an order of magnitude higher delay violations than their actual tolerance. This is because pure over-subscription does not correlate delay violation bound requirements for a QVPN with its bandwidth reservation. We need more than just bandwidth over-subscription – specifically a delay-probability-bandwidth correlation function. such as in Equation 2 – to guarantee distinct per-QVPN probabilistic guarantees.

6.3 Validating the CDF Estimation Technique

Next we validate that the technique for predicting the future delay ratio distribution CDF_{est} in Section 5 indeed reliably bounds the delay ratio distribution of admitted QVPNs. Validating the CDF estimation technique is important to establish that the DDM algorithm does not under-estimate the resource requirements for individual QVPNs, resulting in excess delay violations in the long-term. Figure 8 shows a representative simulation scenario in which 19 constituent QVPNs are admitted with heterogeneous delay bound, data rate, and burstiness requirements. The right-most curve marked CDF_{est} shows the delay ratio distribution estimated by DDM before admitting the 19th QVPN whereas the curves on the left represent the stable per-QVPN distributions at the end of simulation lifetime. The figure demonstrates the fact that the CDF_{est} distribution used at admission control time still remains more conservative than individual QVPN distributions in the long-term. Thus CDF estimation technique can effectively reduce each QVPN's resource requirement to suit their individual tolerance to delay violations without risking under-estimation of true requirements.

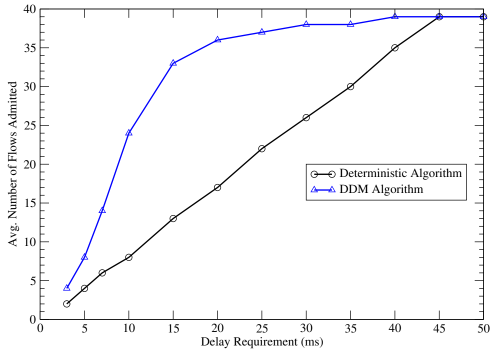

Figure 9: Number of admitted QVPN vs. delay bound. Delay violation probability = 10^{-5}. Burst Size=10pkts. Link capacity = 10Mbps.

6.4 Delay Bound Variation

Next we compare the performance of DDM algorithm against deterministic admission control as delay bound requirement varies. With DDM algorithm, the delay violation probability for each QVPN is 10^{-5} whereas deterministic admission control considers a zero delay violation probability. Figure 9 plots the number of QVPNs admitted as the delay-bound requirement is varied from 3 to 50ms. The maximum number of QVPNs that can be admitted on the 10Mbps link is limited to 39 QVPNs, being limited by the average rate requirement of 256Kbps for each QVPN. Figure 9 shows that for small delay bound requirements, DDM algorithm admits around 3.0 times more number of QVPNs than deterministic admission control when delay violation probability as small as 10^{-5} is allowed. As the delay bound requirement becomes less stringent, DDM algorithm still admits more QVPNs and achieves better link utilization than deterministic algorithm, but with smaller improvements. Beyond 45ms delay requirement, both algorithms are limited to admitting 39 QVPNs due to the average rate requirement of 256Kbps for each QVPN. The gain for DDM algorithm comes from the fact that large majority of packets experience just 1% to 3% of the worst-case delay dictated by their reserved bandwidth. This statistic gets reflected in the CDF which in turn helps to reduce the resource requirement for each QVPN.

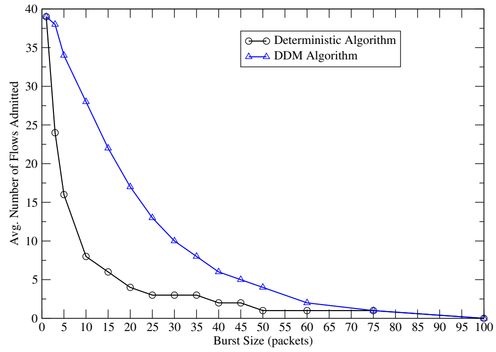

Figure 10: **Number of admitted QVPNs vs. burst size. Delay bound=10ms. Violation Probability=10^{-5}. Link Capacity = 10Mbps.**

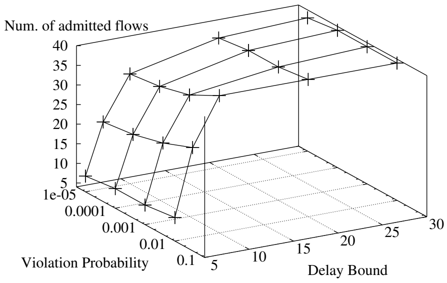

Figure 11: Admission region for various combinations of delay and delay violation probability. Link Capacity = 10Mbps. Burst size=10pkts. Average rate=256Kbps.

6.5 Burst Size Variation

Figure 10 compares the DDM algorithm against deterministic admission control as the burst size σ_i for each QVPN is increased from 1 to 100 packets. Up to burst sizes of 40 packets, DDM algorithm admits significantly larger number of QVPNs than deterministic algorithm since it can successfully exploit the statistical multiplexing among bursts from different QVPNs. For larger burst sizes, the delay-derived bandwidth requirement turns out too high to be adequately compensated by statistical multiplexing.

6.6 Admission Region

Figure 11 shows the admission region for various combinations of delay and delay violation probability. As the delay bound and delay violation probability requirements become less stringent, the number of admitted QVPNs increases. Note that even with a low violation probability of 10^{-5} at 10ms delay, the DDM algorithm can admit up to 24 QVPNs which is 3 times more than deterministic case of 8 QVPNs.

6.7 Effect of CDF Measurement Window

Another factor influencing the performance of the DDM algorithm is CDF measurement window. Figure 12 shows that a large measurement window leads to a more conservative admission process i.e., a large measurement window admits fewer QVPN requests than a small window over the same interval of time. The reason for this behavior can be traced back to Figure 1. Typically bursts from different QVPNs tend to be temporally spread out and multiple

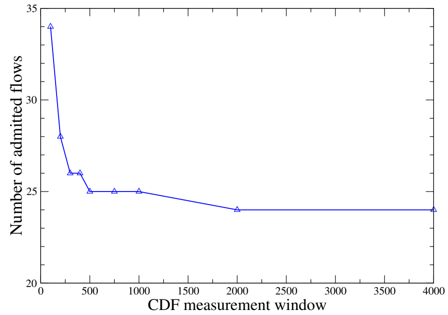

Figure 12: Number of admitted QVPNs with different CDF measurement windows. Admission control becomes more conservative with larger measurement windows. Delay bound=10ms. Violation probability = 10^{-5}. Burst size=10 pkts.

QVPNs rarely burst simultaneously. However, such events do occur and small window sizes are more likely to miss out such rare simultaneous traffic bursts whereas large window sizes are more likely to capture these. Consequently, larger measurement windows produce more representative CDF curves than small windows.

Admission decisions based on small measurement windows could thus be over-optimistic leading to more number of QVPNs being admitted quickly. With large window sizes, the DDM algorithm is slower in reacting to changes in traffic patterns and thus admits fewer QVPNs as traffic load increases. hile a very small window size can result in over-optimistic admissions, an extremely large window size would also lead to inaccurate admission decisions since it might include history that could be too old for consideration. Thus one needs to strike a right balance in selecting a measurement window size that yields optimal performance. A possible choice for the CDF measurement window could be the duration between successive QVPN arrivals since the traffic during this period can be expected to be largely stable and indicative of true load imposed by currently active QVPNs.

6.8 Statistical Multiplexing Gain from Under-utilization

Finally, we vary the number of streams per QVPN to determine the extent of gain we obtain by under-utilizing the aggregate QVPN's reserved capacity. At full capacity, each aggregate QVPN can carry 20 VoIP streams. Figure 13 shows that the number of admitted QVPNs decreases from 35 to 24 as the level of aggregation in each QVPN increases from 2 to 20 VoIP streams. Thus, DDM algorithm can successfully exploit additional statistical multiplexing due to smaller level of aggregation in each QVPN. In this case, the maximum gain is limited by the average rate requirement of 256 Kbps for each QVPN and link capacity of 10Mbps. This is because DDM algorithm exploits the statistical multiplexing effect only along the delay dimension but not along the bandwidth dimension. Multiplexing gains could be higher if the latter dimension could also be accounted in the DDM algorithm.

7. CONCLUSIONS

In this paper, we have proposed a practical measurement-based link-level approach, called *Delay Distribution Mea-*

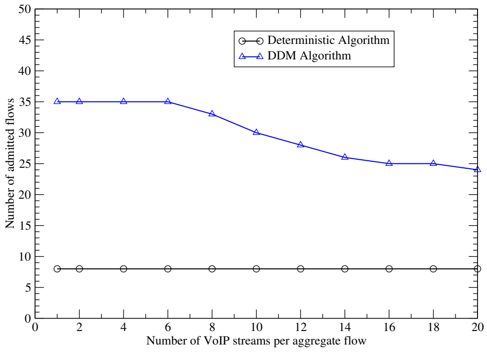

Figure 13: Number of admitted QVPNs with variation in number of VoIP streams per aggregate QVPN. Delay bound=10ms. Violation probability = 10^{-5}. Burst size=10 pkts.

surement (DDM) based admission control, that exploits statistical multiplexing along the delay dimension while providing each QVPN with a distinct probabilistic delay guarantee, i.e. a bound on both delay as well as delay violation probability. By dynamically measuring the distribution of the ratios between actual packet delay and worst-case delay bound, DDM is able to significantly lower the resource requirement of QVPNs that have a small tolerance to delay violations. DDM also provides a unified framework to support QVPNs requiring deterministic or probabilistic delay bounds. Our results, using real VoIP traces, show that the algorithm satisfies heterogeneous probabilistic delay requirements of multiple QVPNs and provides up to 3 times improvement in number of admitted QVPNs even when tolerance to delay violations is as small as 10^{-5}.

The framework of DDM algorithm could be extended to include simultaneous multiplexing along bandwidth dimension as well to yield potentially larger link utilization. We are also interested in using DDM as a building block to exploit statistical multiplexing in the end-to-end scenario where QVPNs traverse multiple network links. Another interesting application of DDM is managing the heterogeneous resources of shared server platforms, such as storage farms or application hosting clusters, where each subscriber receives distinct service rate, latency and tolerance guarantees.

Acknowledgments

We would like to thank Henning Schulzrinne and Wenyu Jiang for providing the VoIP traces and the anonymous referees for their constructive comments.

8. REFERENCES

[1] M. Andrews. Probabilistic end-to-end delay bounds for earliest deadline first scheduling. In *Proc. of IEEE INFOCOM 2000*, March 2000.

[2] R. Boorstyn, A. Burchard, J. Leibeherr, and C. Oottamakorn. Statistical service assurances for traffic scheduling algorithms. *IEEE Journal on Selected Areas in Communications*, 18(13):2651–2664, December 2000.

[3] L. Breslau, S. Jamin, and S. Shenker. Comments on performance of measurement-based admission control algorithms. In *Proc. of IEEE INFOCOM 2000*, March 2000.

[4] C-S. Chang, Y. Chiu, and W. Song. On the performance of multiplexing independent regulated inputs. In *ACM Sigmetrics 2001/Performance 2001*, pages 184–193, 2001.

[5] S. Crosby, I. Leslie, B. McGurk, J. Lewis, R. Russell, and F. Toomey. Statistical properties of a near-optimal measurement-based admission CAC algorithm. In *Proc. of IEEE ATM'97*, June 1997.

[6] R. Cruz. A calculus for network delay, Part I: Network elements in isolation. *IEEE Transactions on Information Theory*, 37(1):114–131, Jan. 1991.

[7] A. Elwalid and D. Mitra. Design of generalized processor sharing schedulers which statistically multiplex heterogeneous QoS classes. In *Proc. of IEEE INFOCOM'99*, pages 1220–1230, March 1999.

[8] S. Floyd. *Comments on measurement-based admission control for controlled load services.* Technical Report, Lawrence Berkeley Laboratory, July 1996.

[9] R. Gibbens and F. Kelly. Measurement-based connection admission control. In *Proc. of 15th Intl. Teletraffic Conference*, June 1997.

[10] K. Gopalan, T. Chiueh, and Y. Lin. Delay budget partitioning to maximize network resource usage efficiency. Proc. of IEEE INFOCOM 2004, Hong Kong, China, March 2004.

[11] F. M. Guillemin, N. Likhanov, R. R. Mazumdar, and C. Rosenberg. Extremal traffic and bounds for the mean delay of multiplexed regulated traffic streams. In *Proc. of IEEE INFOCOM'02, New York, NY*, June 2002.

[12] S. Jamin, P. Danzig, S. Shenker, and L. Zhang. A measurement-based admission control algorithm for integrated services packet networks. *IEEE/ACM Transactions on Networking*, 5(1):56–70, Feb. 1997.

[13] W. Jiang and H. Schulzrinne. Analysis of On-Off patterns in VoIP and their effect on voice traffic aggregation. In *Proc. of ICCCN 2000*, March 1996.

[14] F. Kelly. Notes on effective bandwidths. *In Stochastic Networks: Theory and Applications*, 4:141–168, 1996.

[15] G. Kesidis and T. Konstantopoulos. Worst-case performance of a buffer with independent shaped arrival processes. *IEEE Communication Letters*, 4(1):26–28, Jan. 2000.

[16] E. Knightly and N. B. Shroff. Admission control for statistical QoS. *IEEE Network*, 13(2):20–29, March 1999.

[17] J. Kurose. On computing per-session performance bounds in high-speed multi-hop computer networks. In *Proc. of ACM Sigmetrics'92*, pages 128–139, 1992.

[18] J.-Y. L. Boudec and M. Vojnovic. Stochastic analysis of some expedited forwarding networks. In *IEEE Infocom'02, New York, NY*, June 2002.

[19] A. Parekh and R. Gallager. A generalized processor sharing approach to flow control in integrated services networks: The single-node case. *IEEE/ACM Transactions on Networking*, 1(3):344–357, June 1993.

[20] J. Qiu and E. Knightly. Measurement-based admission control with aggregate traffic envelopes. *IEEE/ACM Transactions on Networking*, 9(2):199–210, April 2001.

[21] M. Reisslein, K. Ross, and S. Rajagopal. A framework for guaranteeing statistical QoS. *IEEE/ACM Transactions on Networking*, 10(1):27–42, February 2002.

[22] V. Sivaraman and F. Chiussi. Providing end-to-end statistical delay guarantees with earliest deadline first scheduling and per-hop traffic shaping. In *Proc. of IEEE INFOCOM 2000*, March 2000.

[23] B. Urgaonkar, P. Shenoy, and T. Roscoe. Resource overbooking and application profiling in shared hosting platforms. In *Proc. of Symposium on Operating Systems Design and Implementation, Boston, MA*, Dec. 2002.

[24] M. Vojnovic and J. Le Boudec. Bounds for independent regulated inputs multiplexed in a service curve network element. *IEEE Transactions on Communications*, 51(5), Jan. 1991.

[25] Y. Wang and Q. Zhu. Error control and concealment for video communication: A review. *Proceedings of IEEE*, 86(5):974–997, May 1998.

[26] L. Zhang. Virtual Clock: A new traffic control algorithm for packet-switched networks. *ACM Transactions on Computer Systems*, 9(2):101–124, May 1991.