University of Pennsylvania
ScholarlyCommons

Departmental Papers (ESE)

Department of Electrical & Systems Engineering

April 2005

Individual QoS versus aggregate QoS: A loss performance study

Ying Xu
University of Pennsylvania

Roch A. Guérin
University of Pennsylvania, guerin@acm.org

Follow this and additional works at: http://repository.upenn.edu/ese_papers

Recommended Citation

This paper is posted at ScholarlyCommons. http://repository.upenn.edu/ese_papers/105
For more information, please contact repository@pobox.upenn.edu.

Individual QoS versus aggregate QoS: A loss performance study

Abstract

This paper explores the differences that can exist between individual and aggregate loss guarantees in an environment where guarantees are only provided at an aggregate level. The focus is on understanding which traffic parameters are responsible for inducing possible deviations and to what extent. In addition, we seek to evaluate the level of additional resources, e.g., bandwidth or buffer, required to ensure that *all* individual loss measures remain below their desired target. This paper's contributions are in developing analytical models that enable the evaluation of individual loss probabilities in settings where only aggregate losses are controlled, and in identifying traffic parameters that have a major influence on the differences between individual and aggregate losses. The latter allows us to further construct tools and guidelines that are able to determine what kind of traffic can be safely multiplexed in practice into a common service class.

Individual QoS Versus Aggregate QoS: A Loss Performance Study

Ying Xu, *Student Member, IEEE,* and Roch Guérin, *Fellow, IEEE*

Abstract—This paper explores the differences that can exist between individual and aggregate loss guarantees in an environment where guarantees are only provided at the aggregate level. The focus is on understanding which traffic parameters are responsible for inducing possible deviations and to what extent. In addition, we seek to evaluate the level of additional resources, e.g., bandwidth or buffer, required to ensure that *all* individual loss measures remain below their desired target. This paper's contributions are in developing analytical models that enable the evaluation of individual loss probabilities in settings where only aggregate losses are controlled, and in identifying traffic parameters that have a major influence on the differences between individual and aggregate losses. The latter allows us to further construct practical tools and guidelines for rapidly assessing if specific traffic sources can be safely multiplexed into a common service class.

Index Terms—Aggregation, loss, quality of service.

I. INTRODUCTION

THE provision of Quality of Service (QoS) guarantees is by now an extensively investigated and reasonably well understood topic. The literature abounds with algorithms for enforcing different levels of services and results evaluating their respective performance, see, e.g., [10] for a recent survey. Similarly, technology is now available that implements sophisticated QoS capabilities, see, e.g., [15]. However, despite all this progress, the *deployment* of QoS capabilities in operational networks has been by most accounts *slow*. Many factors have conspired toward this, but one particular factor has been a recurring theme in discussions aimed at understanding the reasons behind this slow pace. Specifically, the complexity of managing a broad range of fine grain (individual) QoS requirements across a network of the scale of the Internet is a daunting task. As a result, there has been a renewed interest in designing *scalable* QoS solutions.

There have been two main directions aimed at developing scalable QoS solutions. The first, embodied in works such as [23], [24], [28], targets the emulation of fine-grain QoS solutions without requiring per flow information. The second, represented by proposals such as Diff-Serv [4], relies on coarsening the different levels of QoS that the network offers into a small number of service classes. Our focus is on this latter class of solutions.

Limiting the number of service classes that the network offers clearly improves scalability. However, this comes at a cost, namely the lack of awareness of the exact level of performance that an individual user experiences. In other words, implicit in the use of service classes is the assumption that the users aggregated into a given service class all experience the *same* level of service, or at least a level of service better than the desired target for the service class. Unfortunately, this assumption may not always be true in an environment where *provisioning* and class level *monitoring* are the primary tools used to enforce network performance, and make decisions on whether additional traffic can be accommodated.

By its nature and in order to ensure its scalability, provisioning is typically done at the aggregate level, e.g., based on the monitored performance of a service class. In that context, our objective in this paper is to gain a better understanding of how such aggregate measures map into individual performance, and in particular determine if guidelines and procedures could be developed to avoid situations where the two differ significantly. In contrast to environments where call admission procedures, e.g, [7]–[9], [13], are used to dynamically make decisions on accepting new traffic, we focus on settings where these guidelines and procedures are applied for making *off-line* decisions about whether to multiplex different types of traffic into a common service class.

In our study, we focus on a specific performance measure, namely, the packet loss probability. To gain a comprehensive understanding of how individual loss probabilities differ from overall loss probability, we first develop a number of new models or extensions to existing models, which allow us to analytically evaluate the loss probability experienced by an individual flow when only the overall loss probability of the service class to which this flow belongs is observable. With a broad enough coverage of the entire parameter space, these models and extensions enable us to investigate the influence different traffic parameters, e.g., peak rate, average rate, burst duration, etc., have on introducing deviations between a flow's own loss probability and the overall loss probability. In that context, the identification of parameters and situations that can lead to significant deviations is of special interest. In addition, we also evaluate the sensitivity of performance deviations to the amount of additional resources, i.e., bandwidth or buffer, required to ensure that even the worst performer in a service class experiences a level of performance equal to or better than the desired target for the service class. The result of these

The authors are with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104 USA (e-mail: yingx@ee.upenn.edu; guerin@ee.upenn.edu).

investigations is a set of guidelines and recommendations that identify if and when aggregation can be safely performed.

Moreover, we investigate the extent to which these guidelines and recommendations apply to real-life scenarios. For that purpose, we incorporate those guidelines and recommendations into a simple methodology capable of predicting the presence and magnitude of deviations when multiplexing actual traffic flows. The methodology is still rooted in the analytical models that were developed, but involves several simplifying assumptions for accommodating realistic traffic sources. The performance of this methodology is tested via simulations using a number of real-time traffic sources, i.e., voice and/or video traffic sources. The results show that the methodology is reasonably robust and reliable in identifying situations where significant loss performance deviations can arise. Thus, it can serve as a simple tool for rapidly characterizing traffic profiles that should not be aggregated into a common service class.

The rest of this paper is structured as follows. Section II motivates and introduces the different models and systems assumptions. Analytical expressions of individual loss probabilities are provided for those models in Section III, with most of the derivations relegated to an extended technical report [26]. Sections IV and V report results obtained from the models of Section III, while a number of additional intermediate configurations are investigated in Section VI. In Section VII, the methodology of assessing loss deviations in practical settings is presented and its performance is also explored. Finally, Section VIII summarizes the main findings of the paper and their implications for aggregate QoS solutions.

II. MODEL AND METHODOLOGY

In this section, we describe the model and methodology we rely on to investigate the behavior of individual loss probabilities in an environment based on aggregate service classes. This includes the source traffic model, the different service configurations we analyze, and how we measure differences between individual and aggregate loss probabilities. Specifically, the system we consider is a single server, finite buffer, FIFO queue where traffic generated by users of a common service class is aggregated.

A. Input Traffic Model

In this subsection, we present the traffic models we use as traffic sources. They include both analytical models and traffic traces generated by real applications.

1) Analytical Traffic Sources: We consider two different analytical models for characterizing the traffic of an individual source feeding the FIFO queue. The first is a standard ON-OFF Markov source [2], with exponentially distributed ON and OFF periods and a fixed transmission rate when ON (active). Such a source can be described using a 3-tuple $\langle R, b, \rho \rangle$, where R is the transmission (peak) rate when the source is active, b is the average duration of an active or ON period, and ρ represents the

fraction of time the source is active, or its utilization. The rationale for such a source model is both that it lends itself to the development of tractable analytical models from which intuition and insight can be derived, and that its simple three-parameters description can be easily mapped onto popular traffic control devices such as leaky buckets, e.g., see [12] for a discussion on this issue. As a result, it captures the behavior of configurations where performance is mainly determined by "burst-level" congestion. This will be the case when provisioning used for the service class allows for periods of time during which the incoming traffic rate exceeds the allocated capacity.

The second model considered is the $\sum D_i / D / 1$ queue [19], [20], [25]. In this model, each source has a constant bit rate and periodically (every D_i units of time) generates a single, unit size[1] packet. Sources can differ in terms of both their periods and phases, i.e., the positions at which they generate their packets in a period. The choice of the phase of a source is assumed to be independent of that of other sources, and to be drawn from a uniform distribution over the interval $[0, D_i)$. In contrast to the first model, this second model captures an environment in which congestion primarily occurs at the packet level. This is the case when provisioning is done based on the "worst-case" assumptions regarding the traffic that a user can generate. For example, this could apply to provisioning rules used to support a constant rate, low delay service class based on the Diff-Serv Expedited Forwarding Per Hop Behavior (PHB) [5], or representative of a network that relies on conservative (peak rate) provisioning.

2) Real-World Traffic Sources: In addition to the analytical traffic models, we also use real world traffic sources for our study. We mainly focus on real-time traffic sources such as voice and/or video sources, as they correspond to applications with a need for service guarantees, and therefore, most likely to benefit from QoS solutions. From that perspective, we believe that evaluating performance deviations in environments where voice and video sources are aggregated can also help us to gain a better understanding of if and when an aggregate service model is suitable for supporting real-time applications.

The video traces we use for our study are obtained from an online video trace library [1]. We choose this public library because of the diversity of the traces it provides. The library contains traces of 26 different hour-long video sequences, each encoded using two different standards (the MPEG4 and the H.263 standard), and under three different quality indexes (low, medium, and high). Moreover, for each recorded trace, statistics reflecting its characteristics are also available. We refer to [6] for a detailed description of the trace collection procedure.

We mainly focus on two different sequences encoded using the MPEG4 encoding technique with a frame rate of 25 frames/s. They include both the high and low quality sequences of the movie *Jurassic Park I*. The high quality trace of *Jurassic Park I* (Jurassic high) represents a video source that requires a high transmission rate due to its quality requirement. The low quality trace of *Jurassic Park I* (Jurassic low) represents a lower rate sequence encoded from the same movie scenes. Several important statistics of these two traces are given in Table I.

---

[1]Its transmission takes one unit of time.

TABLE I
VIDEO TRACE STATISTICS

Trace Name	Jurassic low	Jurassic high
Mean Frame Size (Bytes)	768.6	3831.5
Mean Bit Rate (Kb/sec)	153.7	766.3
Frame Size (Peak/Mean)	10.6	4.4
Frame Size (Covariance)	1.4	0.6

TABLE II
VOICE TRACE STATISTICS

Trace Name	Voice
Mean Spurt Length (ms)	326
Mean Gap Length (ms)	442
Peak Rate (Kb/sec)	64
Utilization	0.425

We also use voice traces published in [14] for our study. In particular, the voice source[2] we chose corresponds to [14, Fig. 4b] and is generated by the NeVoT Silence Detector (SD) [21]. We refer to [14] for a complete description of the codec configuration and the trace collection procedure. Statistics for several important parameters of the voice source are given in Table II.

B. System Parameters and Performance Measures

In developing models to explore possible deviations between individual and aggregate loss probabilities, we vary several system parameters to sample a comprehensive range of possible environments.

The first parameter we vary is the number of users aggregated in the same service class, i.e., the number of individual traffic sources multiplexed in the FIFO queue. In particular, we focus on two cases: a two-source configuration and a "many-source" one. We select these two configurations not only because they are amenable to analysis, but also because they correspond to different boundary conditions, i.e., an environment where a few large bandwidth connections share resources, and one where many small (compared to the link capacity) flows are multiplexed into the same queue. We expect these two environments to exhibit different sensitivity to the traffic parameters of individual sources, and to possibly yield deviations of different magnitude.

Another system parameter we consider is the size of the FIFO queue into which flows are multiplexed. In particular, we consider both bufferless (buffer size of zero) and buffered systems. In many instances, bufferless systems lend themselves to a more tractable analysis, while qualitatively capturing performance trends. In cases where models are limited to bufferless systems, we also rely on simulations to extend the investigation to buffered systems. In all cases, the simulation

results confirmed the trends observed from the bufferless analytical results.

Our focus is to derive explicit expressions, functions of source and system parameters, for *both* individual and aggregate loss probabilities for the different system configurations we consider. The loss probability is computed as the ratio of the total number of bits lost to the total number of bits sent, for either an individual user or all users. We denote the overall loss probability as P_L, and the loss probability of user i as P_L^i. Note that traditional models, e.g., [2], [22], have focused on deriving expressions for the (overall) *overflow* probability rather than the loss probability, where the overflow probability denoted the fraction of time that the system was losing data. The latter, however, is in our minds a more realistic measure of the performance that individual users experience. As outlined in Section III, deriving expressions for the loss rather than the overflow probability calls for some slight modifications to the models developed in [2] and [22].

The availability of expressions for individual and aggregate loss probabilities allows us to investigate if and when they differ. We carry out these investigations by first selecting a target loss probability $P_{\max}$ for the service class, and then computing the (*minimum*) amount of bandwidth C needed to ensure an overall loss probability $P_L \leq P_{\max}$. For each individual user, its loss probability ratio PLR_i is defined as the ratio between its own loss probability and the overall loss probability. We then compute the maximum loss probability ratio across all users, PLR, and select it as the basic performance measure for evaluating the level of deviations in a given configuration. In a system with N users, the user having the maximum loss probability ratio is, without loss of generality, usually assumed to be user N, i.e., $\mathrm{PLR} = P_L^N / P_L$. In addition, we also evaluate the amount of additional bandwidth needed, percentage wise, to ensure that even the worst performer in the service class meets the desired loss target. In other words, if C_N denotes the minimum amount of bandwidth needed so that $P_L^N \leq P_{\max}$, we evaluate the quantity $(C_N - C)/C$. Note that the (maximum) loss probability ratio reflects the magnitude of the loss performance deviations and is a useful metric for determining whether traffic aggregation should be carried out. Conversely, the additional bandwidth is a useful metric for evaluating the amount of additional resources needed to compensate for any deviation, a penalty that an aggregate service model incurs.

We believe that the combined use of analytical models and actual traffic sources together with the investigation of a broad range of system configurations, allows for a reasonably comprehensive exploration of the problem space.

III. ANALYTICAL MODELS

This section is devoted to the presentation of the analytical models that allow us to compute and compare both individual and aggregate loss probabilities. Due to space constraints, proofs and additional details are relegated to a technical report [26]. A total of four distinct models are described in this section. The first two correspond to two-source cases assuming first ON-OFF and then periodic sources; the last two correspond

---

[2]We use only one voice trace since in our study, voice sources are mainly used to evaluate the presence of deviations when mixed with video sources. Because of their statistical similarities, different voice sources experience similar losses when aggregated.

to the "many-source" scenario, again for ON-OFF and periodic sources.

A. Two ON-OFF Sources Case

Since there is conceptually little difference in the derivation of expressions for a two-source system and an N-source system, we proceed to derive general expressions for an N-source system, which is then specialized to a two-source system by letting $N = 2$. The analysis of systems that involve Markov modulated traffic sources is by now a mature area, e.g., [2], [17], [22], and we rely on this existing body of work to develop our model. The main differences between those works and ours are that we focus on the loss probability as opposed to the overflow probability and, most important, we evaluate both individual and aggregate performance.

Our initial model consists of N independent ON-OFF fluid sources that feed an infinite buffer, single server queue. Source i is characterized by a 3-tuple (R_i, b_i, ρ_i) as described in Section II-A1. The aggregate input process to the buffer can then be described through a state vector: $\mathbf{S} = (s_1, s_2, \ldots, s_N)$, where s_i is 0 when source i is OFF and 1 when it is ON. For any state, the input rate γ_S to the system is given by $\gamma_{\mathbf{S}} \doteq \mathbf{S} \cdot \mathbf{R^T}$, where $\mathbf{R} = (R_1, R_2, \ldots, R_N)$ is the peak rate vector of the sources. Let $\pi_{\mathbf{S}}$ denote the stationary probability that the input is in state $\mathbf{S}$, under the standard assumption that the system is ergodic, the stationary loss probability experienced by source i in a finite buffer system of size x can be approximated by

$$P_L^i = \frac{\sum\limits_{\mathbf{S}:(s_i=1,\gamma_{\mathbf{S}}>C)} \left(\pi_{\mathbf{S}} - F_{\mathbf{S}}(x)\right)\left(R_i - C \cdot \frac{R_i}{\gamma_{\mathbf{S}}}\right)}{\rho_i R_i} = \frac{r_L^i}{r_S^i} \quad (1)$$

where r_L^i and r_S^i correspond to the long term loss rate[3] and sending rate (mean rate) of source i. The quantity $F_{\mathbf{S}}(x)$ is the stationary probability that the queue length is smaller than x and the system is in state $\mathbf{S}$, which can be readily obtained from results of either [17] or [22].

Similarly, the overall loss probability P_L can be expressed as

$$P_L = \frac{\sum\limits_{\mathbf{S}:\gamma_{\mathbf{S}}>C} \left(\pi_{\mathbf{S}} - F_{\mathbf{S}}(x)\right)(\gamma_{\mathbf{S}} - C)}{\sum\limits_{i}^{N} \rho_i R_i} = \frac{r_L}{r_S} \quad (2)$$

where r_L and r_S correspond to the overall long term loss rate and sending rate, respectively.

Expressions for individual and aggregate loss probabilities can be readily obtained from (1) and (2) for the two-source case simply by letting $N = 2$. For example, in the case where $R_1 \leq C$, $R_2 \leq C$, and $R_1 + R_2 > C$, i.e., losses occur only when both sources are active, we have

$$P_L^1 = \frac{\left(\pi_{(1,1)} - F_{(1,1)}(x)\right) \cdot (R_1 + R_2 - C)}{\rho_1(R_1 + R_2)}$$

$$P_L^2 = \frac{\left(\pi_{(1,1)} - F_{(1,1)}(x)\right) \cdot (R_1 + R_2 - C)}{\rho_2(R_1 + R_2)} \quad (3)$$

[3]We use the notion of "loss rate" to distinguish the amount of data lost per unit of time from the fraction of data lost. The latter is termed loss probability in this paper.

where $\pi_{(1,1)}$ is the stationary probability that both sources are active. As we shall discuss further in Section IV, the simple form of (3) helps explicitly identify the impact of different parameters. In particular, we see that for this special case, the ratio of individual loss probabilities is equal to the inverse ratio of the utilizations of the respective sources. In other words, the source with the lower utilization will see a proportionally higher loss probability. This simple but nevertheless interesting observation is one that will be subsequently confirmed in other and more general configurations.

B. Two Periodic Sources Case

In this section, we consider one special case of the $\sum D_i/D/1$ queue in which there are only two sources in the system, with periods D_1 and D_2, respectively, i.e., the $(D_1 + D_2)/D/1$ queue. For this model, we assume a bufferless system, since having a buffer of size larger than or equal to even *one* packet will eliminate all losses. In spite of its extreme simplicity, this system is again useful because of the insight it provides. For this simple system, the following proposition can be shown to hold.

Proposition 1: For a *bufferless* $(D_1 + D_2)/D/1$ queue, where D_1, D_2 are *integers* and $2 \leq D_1 \leq D_2$, the loss probability ratios P_L^2/P_L^1 and P_L^2/P_L satisfy

$$\frac{P_L^2}{P_L^1} = \frac{D_2}{D_1} \quad \frac{P_L^2}{P_L} = \frac{1 + \frac{D_2}{D_1}}{2}. \quad (4)$$

Moreover, if D_1 is divisible by D_2, then

$$P_L^1 = \frac{1}{D_2} \quad P_L^2 = \frac{1}{D_1} \quad P_L = \frac{2}{D_1 + D_2}. \quad (5)$$

Proof: The proof is given in [26]. ∎

The above proposition, though simple, states an important fact that the source with a larger period sees greater losses in proportion to the ratio of its period over that of the other source. This is an observation similar to the one made based on (3) for two ON-OFF sources, and it will again occur in several other configurations, as we shall see later.

C. Many ON-OFF Sources Case

Equations (1) and (2) were derived for the general case of N sources, but as N grows large, the required computations quickly become prohibitive because of the well-known "state explosion" phenomenon, and this makes numerical evaluation difficult if not impossible. As a result, we rely on simulations for evaluating buffered systems, and shift the focus of our analysis to a bufferless system for which numerical evaluation remains feasible. A bufferless model is a reasonable approximation when a large number of sources are multiplexed. In this case, the input traffic can be well approximated using the technique of rate envelope multiplexing [19, sec. 4.1.1]. The rate envelop of a traffic aggregate is simply the instantaneous rate of the total multiplexed traffic. With this notion, data loss can be characterized by relying on the fact that it only occurs when the input rate (envelope) exceeds the available service rate and a small loss probability can always be achieved by ensuring that this event occurs with a sufficiently low probability. Under these

assumptions, models already exist [18] that give explicit expressions for the quantities of interest in the context of this paper, i.e., individual and aggregate loss probabilities. For completeness, we briefly restate the relevant results and assumptions of [18]. Note that [18] is one of the first works to explicitly target understanding when and why differences in performance can arise when aggregating many different types of sources. Many of its results are consistent with those we derive in this paper, and the main differences are both in terms of the broader investigation we undertake, and more important, of our focus on explicitly identifying the impact of individual traffic parameters on performance deviations.

In the bufferless system we consider, the total input traffic is divided into two parts: the background traffic and the traffic associated with a specific source. This source (without loss of generality, we assume it is source N) is the one we focus on, and whose traffic parameters we vary. We denote by λ_t, λ_t^b, and λ_t^N as random variables associated with the instantaneous rate (envelop) of the total traffic, the background traffic, and the traffic generated by source N at time t, respectively. Similarly, the variables m, m_b, and m_N identify the corresponding mean rates. The overall loss probability P_L and the loss probability P_L^N of source N can then be obtained through a minor generalization of the results of [18], by assuming a link capacity of C instead of a unit link capacity

$$P_L^N = \frac{E\left[(\lambda_t - C)^+ \cdot \frac{\lambda_t^N}{\lambda_t}\right]}{\rho_N * R_N}$$
$$= E\left[\left(\frac{\lambda_t^b + R_N - C}{\lambda_t^b + R_N}\right)^+\right] \qquad (6)$$

and

$$P_L = \frac{E\left[(\lambda_t - C)^+\right]}{m}$$
$$= \frac{\rho_N E\left[\left(\lambda_t^b + R_N - C\right)^+\right] + (1 - \rho_N) E\left[\left(\lambda_t^b - C\right)^+\right]}{m}. \qquad (7)$$

The expectations in both (6) and (7) can be evaluated numerically if λ_t^b is explicitly specified. In most of our tests, the background traffic consists of homogeneous sources with identical peak rate R_b, as we focus on the impact of varying the traffic parameters of source N. In such cases, λ_t^b / R_b is simply a binomial distribution.

D. Many Periodic Sources Case

The overflow probability of the $\sum D_i / D / 1$ queue has been investigated in [19], [20], and [25]. Its derivation is based on the Beneš approach that is extensively documented in [19]. In this paper, we build on and extend the methods of [20] and [25] to obtain upper and lower bounds for individual loss probabilities. We briefly outline the model assumptions and state the final results, while details and proofs can be found in [26].

The model we consider consists of J different types of sources. There are N_j independent sources of type j, $1 \leq j \leq J$, each with a period of D_j. The individual source we focus on is source i with period D_i. In order to

ensure that the system is stationary and ergodic, we further impose the condition that the total load ρ is less than 1, i.e., $\rho = 1/D_i + \sum_{j=1}^{J}(N_j/D_j) < 1$. The queue length, or the backlog, in the system at time t is denoted as V_t, and we use the expression

$$\Pr\{V_0 > x | \text{One arrival from source } D_i \text{ at } 0^+\}$$

to approximate the loss probability of source D_i in a finite buffer system of size x. In [26], the following lower and upper bounds are established:

$$P_L^i \geq \sum_{\substack{n > x \\ n \neq x + \lceil ND_i \rceil - 1}} (1 - \rho)$$
$$\cdot \left(\sum_{K(n,x) = n - d(n,x)} \prod_{j=1}^{j} q_j\left(k_j(n,x)\right) \right) \qquad (8)$$

$$P_L^i \leq \sum_{\substack{n > x \\ n \neq x + \lceil ND_i \rceil - 1}}$$
$$\cdot \left\{ \sum_{K(n,x) = n - d(n,x)} \left[\prod_{j=1}^{j} q_j\left(k_j(n,x)\right. \right. \right.$$
$$\left. \left. \left. \cdot \left(1 - \frac{N_j - k_j(n,x)}{D_j\left(1 - p_j(n,x)\right)}\right)^+ \right] \right\} \qquad (9)$$

where $d(n,x) = \sum_{j=1}^{J} N_j \cdot \lfloor (n-x)/D_j \rfloor + \lfloor (n-x)/D_i \rfloor$, $p_j(n,x) = ((n-x)/D_j) - \lfloor (n-x)/D_j \rfloor$, and $K(n,x) = \sum_{j=1}^{J} k_j(n,x)$. The value of $k_j(n,x)$ can vary between 0 and N_j and $q_j(k_j(n,x))$ is the probability density function of a binomial distribution with parameter $(N_j, p_j(n,x))$. See again [26] for more complete definitions.

Upper and lower bounds for the overall loss probability P_L can also be obtained from (9) and (8) by bringing them into the following expression:

$$P_L = \frac{\sum_{j=1}^{J} \frac{N_j}{D_j} \cdot P_L^j + \frac{1}{D_i} \cdot P_L^i}{\sum_{j=1}^{J} \frac{N_j}{D_j} + \frac{1}{D_i}}. \qquad (10)$$

The above upper and lower bounds have been numerically evaluated for many different configurations, and found to be consistently very close to each other (see [26] for a couple of illustrative examples). In this paper, we use the upper bound (9) to approximate the individual loss probability P_L^i of source i.

In the next sections, we rely on the various expressions derived in this section to investigate when and why individual and aggregate loss probabilities differ. In all our investigations, we assume a fixed target loss probability $P_{\max}$ equal to 10^{-4}.

IV. Loss Deviations in the Two-Source Case

This section is devoted to exploring configurations that involve only two sources.

A. Two ON-OFF Sources

In this section, we assume that only two ON-OFF sources are multiplexed into a common queue served by a constant rate
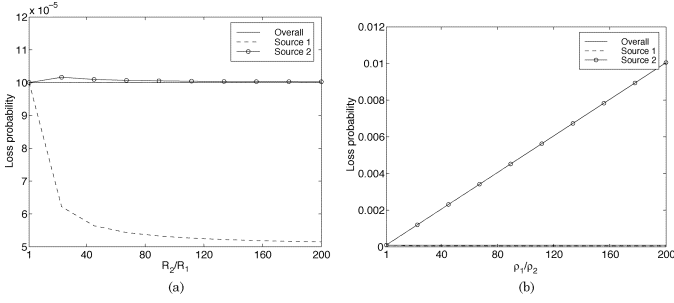
Fig. 1. Performance deviations for two ON-OFF sources. (a) Varying R_2. (b) Varying ρ_2.

server with a speed of C bits/s, where the value of C has been selected to ensure an overall loss probability $P_L = 10^{-4}$. The traffic parameters of source 1 are kept fixed at $R_1 = 10^8$ bits/s, $b_1 = 0.005$ s, and $\rho_1 = 0.5$. In all scenarios described in this section, the buffer size of the queue is set equal to the total average burst size of the two sources. Additional experiments were conducted with different buffer sizes, and did not yield drastically different behaviors. The parameters of source 2 are varied one or more at a time, with its other parameters kept constant and identical to those of source 1. We rely on (1) and (2) to evaluate the deviations between P_L^2 and P_L. We omit results related to varying the burst duration as this parameter was found to have no or only minor impact. This does not mean that the burst duration has no impact. It certainly does. As it increases, the total allocated bandwidth C increases to accommodate the burstier arrival process of source 2. However, contrary to what happens with the peak rate and the utilization, varying the burst duration while maintaining the aggregate loss probability below the desired target of $P_L = 10^{-4}$, does not introduce significant difference between the loss probabilities of two sources.

The first set of conclusions one can draw from this simple configuration, is that individual traffic parameters can indeed induce loss probability differences across sources. This is illustrated in Fig. 1. In particular, we see that either increasing the peak rate of source 2 or decreasing its utilization translates into source 2 experiencing a higher loss probability. However, as illustrated in Fig. 1(a) and (b), the amount by which the loss probability of source 2 exceeds the overall loss probability is quite different in these two cases.

The fact that both peak rate and utilization can affect the performance of an individual source is reasonably intuitive. A higher peak rate source dumps data faster into the buffer, which increases its likelihood of losing data. This increased burstiness notwithstanding, the difference between P_L^2 and P_L remains small, as seen in Fig. 1(a). This is primarily because of how we vary individual traffic parameters, and the resulting weight of each source in terms of its traffic contribution. Specifically, when we vary (increase) the peak rate of source 2, because the utilization of the two sources remains the same, source 2 ends up being the dominant contributor of traffic to the system. Hence, although it does experience higher losses, because of its higher weight in computing the overall loss probability P_L, the allocated bandwidth C that is chosen to ensure that $P_L \leq 10^{-4}$, also ensures that P_L^2 remains close to this target value.

The relationship between the fraction of the traffic contributed by an individual user and the *maximum* ratio between its own loss probability and the overall loss probability can be characterized by a simple upper bound given below. Assuming an environment where there are a total of N users, and following the same set of notations used in Section III-A, the loss probability ratio of user i should satisfy

$$\text{PLR}_i = \frac{P_L^i}{P_L} = \frac{r_L^i/r_S^i}{r_L/r_S} \leq \frac{r_L^i/r_S^i}{r_L^i/r_S} = \frac{1}{r_S^i/r_S} \qquad (11)$$

where the equality holds if and only if source i is the only source losing data. The above equation clearly states that the loss probability ratio experienced by a user can never exceed a value that is inversely proportional to the fraction of traffic it contributes to the total input traffic. For the case of varying the peak rate, the above upper bound indicates that source 2 can never have a loss probability ratio larger than 2, since it always contributes more than half of the total input traffic. Furthermore, as the peak rate of source 2 increases, so will the fraction of traffic it contributes. Hence, the maximum possible value of its loss probability ratio decreases according to (11). For example, when $R_2/R_1 = 200$, from (11) we have $\text{PLR}_2 \leq 1.005$. In short, (11) implies that major traffic contributors can at worst only see minor performance degradations, a theme that we will encounter repeatedly throughout this paper.

In contrast, as is seen in Fig. 1(b), when we only vary the utilization of source 2, the difference between the individual and aggregate loss probabilities can be much larger. This is because when the utilization of source 2 is varied (decreased) while its peak rate remains identical to that of source 1, it is source 1 that becomes the dominant traffic contributor. Hence, the allocated bandwidth C is determined primarily based on the performance of source 1, which allows the loss curve of source 2 to degrade almost arbitrarily [PLR_2 could be as large as 201 in the worst case according to (11)]. The reason for this degradation is that the lower utilization of source 2 limits its ability to access the link. This provides source 1 with additional transmission opportunities, which help lower its individual loss probability. This is best illustrated through the scenario mentioned in Section III-A, where losses occur only when both sources are active.

In this special case, the individual loss probabilities are given in (3), which clearly identifies the impact of the smaller utilization of source 2. Specifically, we have $P_L^2 = P_L^1(\rho_1/\rho_2)$, which increases in a way that is inversely proportional to ρ_2. This is because losses occur only when both sources are active and are, therefore, distributed in proportion to their peak rates, i.e., $r_L^i \propto R_i$. Meanwhile, the transmission rate r_S^i of an individual source, which reflects its transmission opportunities, is proportional to both the peak rate R_i and the utilization ρ_i. Thus, the effect of the peak rates R_i cancels out, and P_L^2/P_L^1 can be easily found to be inversely proportional to ρ_2/ρ_1. A similar trend was also observed in cases where the peak rate of each source exceeds the link capacity. However, there are two competing effects in those cases, namely, how the peak rate of a source influences how much it loses when it is the only one active, and how its utilization affects its ability to gain access to transmission opportunities. From our observations, utilization
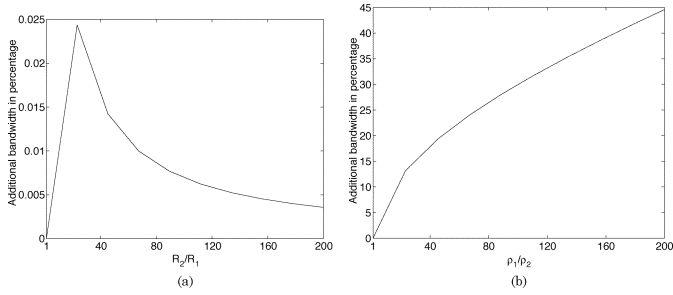
Fig. 2.   Additional bandwidth for two ON-OFF sources. (a) Varying R_2. (b) Varying ρ_2.

Fig. 3.   Loss probability deviations for many ON-OFF sources. (a) Varying R_2. (b) Varying ρ_2. (c) Varying R_2 ρ_2.

remains the dominant factor, primarily because the allocated capacity is typically chosen so as to keep losses at a sufficiently low level when only one source is active.

The next aspect we investigate is the amount of additional resources required to ensure that *both* sources experience a loss probability that is below the desired target of 10^{-4}. The results of this investigation are shown in Fig. 2 in the form of the percentage of additional bandwidth required. In particular, Fig. 2(b) confirms the potentially severe penalty imposed by mixing sources with very different utilizations, as the amount of additional bandwidth needed can reach about 45%.

In [26], we carried out additional experiments, where we simultaneously varied two instead of only one traffic parameters of source 2. In particular, we considered scenarios that involved varying R_2 and b_2, R_2 and ρ_2, and b_2 and ρ_2. The results of changing $R_2\&b_2$ and changing $b_2\&\rho_2$ were quantitatively very similar to the cases of changing R_2 and ρ_2, respectively. This again confirms that the burst size can only have minor impact on inducing loss deviations. When both R_2 and ρ_2 were varied, differences between the two sources were also minor, and a small amount of additional bandwidth was sufficient to guarantee both their performance. Due to space constraints, we refer the reader to [26] for details on these additional scenarios.

B. Two Periodic Sources

From Proposition 1, we know that as for ON-OFF sources, the source with the longer period (smaller rate) experiences higher losses in proportion to the ratio of its period to that of the shorter period (higher rate) source. The reasons are again similar to those articulated for ON-OFF sources. Specifically, a periodic source with a longer period, like an ON-OFF source with a smaller utilization, has fewer opportunities to access the link and transmit a packet, and therefore experiences a higher loss probability.

Moreover, we know from (5) that when D_1, D_2 and their ratio $D_2/D_1 = N$ are all integers, the additional bandwidth required to ensure a target loss probability of $P_{\max}$, when expressed in terms of the corresponding increase of the sources' period, is equal to $(N-1)/2 = (D_2/D_1 - 1)/2$. Note that this value is independent of $P_{\max}$ and can be made arbitrarily large by increasing D_2/D_1.

V. LOSS DEVIATIONS IN THE MANY-SOURCE CASE

This section targets what one can consider a more realistic set of scenarios, namely, service classes that carry a large number of
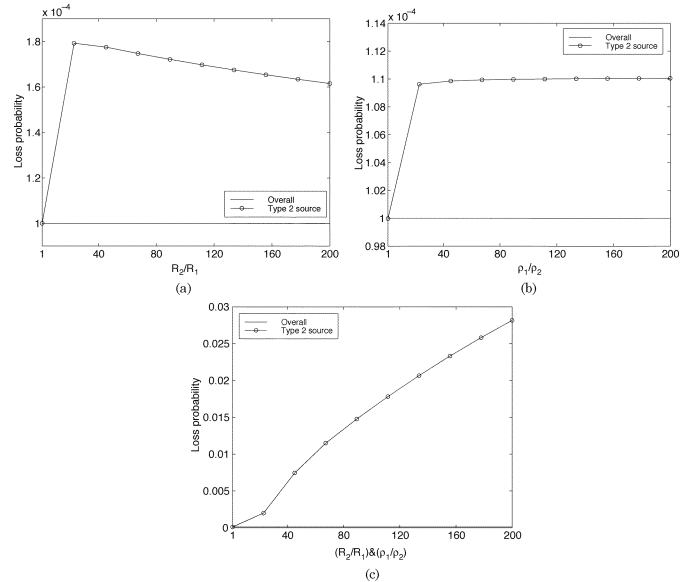
flows. Such configurations will clearly be more appropriate for high speed links, where one can expect to see just a few service classes, e.g., built on top of a small numbers of Diff-Serv PHBs, each carrying a large number of flows. In such an environment, our intuitive expectation is that the presence of a large number of flows is likely to "soften" possible deviations in performance. As we will see, this intuition will indeed be confirmed.

A. Many ON-OFF Sources

In this subsection, we assume that the input traffic consists of many ON-OFF sources. For simplicity, we limit ourselves to only two types of sources. This facilitates the identification of which user (type) experiences the higher losses. The type 1 sources form the "background" traffic, and we assume a total of 1000 such sources, each with a peak rate $R_1 = 10^5$ bits/s, a burst duration $b_1 = 0.005$ s, and a utilization $\rho_1 = 0.5$. As previously mentioned in Section III-C, when the background traffic is homogeneous, the number of active background sources follows a binomial distribution. Using this fact together with (6) and (7), we can evaluate the loss probability experienced by both the type 1 sources and a single type 2 source whose traffic parameters we vary. Since the burst duration has no impact in a bufferless model, we first vary R_2 and ρ_2 one at a time, and then vary both of them while keeping their product constant, i.e., $R_2\rho_2 = R_1\rho_1$. Again, when any parameters of the type 2 source are varied, its other parameters are set to the same values as those of the type 1 sources.

The results are shown in Fig. 3. The figure illustrates that variations in either peak rate or utilization alone do yield some differences between the two types of sources. However, those differences are most significant when the type 2 source has both a much higher peak rate and a much lower utilization than the type 1 sources. In particular, Fig. 3(c) shows that when $R_2/R_1 = \rho_1/\rho_2 = 200$, the loss probability ratio is about 280. Such a result is reasonably intuitive. First, when the peak rate of the type 2 source increases, so does its impact, making it more likely to create congestion when active. The extent to

which this also triggers the allocation of additional bandwidth to compensate for this potential increase in losses depends on the impact of the type 2 source on the overall loss. In particular, if its utilization is very low, such impact will be minor and will not trigger the allocation of any substantial additional bandwidth. Hence, bandwidth allocation is primarily driven by the performance of the type 1 sources, and the higher losses of the type 2 source will remain mostly undetected. This explains why the type 2 source can experience losses much higher than the aggregate loss when both its peak rate and utilization are varied while its overall data rate is kept constant.

In contrast, when only the peak rate of the type 2 source is varied (increased), so will its overall data rate, which means that the weight of its losses in the overall loss computation will also increase. This will typically ensure that those losses are properly accounted for in the bandwidth allocation procedure, so that the bandwidth allocated to keep the aggregate loss probability constant is sufficient to also keep the losses of source 2 close to the desired target. This can be observed in Fig. 3(a), where the loss probability of the type 2 source is only slightly larger than the overall loss probability. The figure also shows that the relative performance of the type 2 source initially degrades as its peak rate increases, and then gradually improves. The presence of such a "cross-over" point is merely a reflection of the fact that there is a lag between the negative impact of a higher peak rate on the performance of the type 2 source and the eventual detection of such impact by the bandwidth allocation procedure as it increases.

Fig. 3(b) illustrates that when only the utilization of the type 2 source is varied, the differences between the loss performance of the two types of sources are marginal. This is because although decreasing the utilization of source 2 also decreases its traffic contribution, and thus its impact on the bandwidth allocation procedure, source 2 only has a minor impact on the overall congestion of the system when it becomes active. Hence, it sees a system nearly identical to what the type 1 sources see, and will therefore experience a similar loss probability.

As for the two-source case, we also investigated the amount of additional bandwidth needed to ensure that the type 2 source experiences the desired target loss. The results are reported in Fig. 4. The figure shows that across all the scenarios of Fig. 3, the maximum amount of additional bandwidth needed is only 8%, as compared to levels in excess of 40% for some of the two-source scenarios. This further confirms our earlier intuition regarding the benefits of larger scale systems toward guaranteeing consistent performance across users. However, note that in terms of absolute value, the additional bandwidth required to adequately accommodate a low utilization, high peak rate source remains high.

Because the above results were obtained using a bufferless model, we performed a similar set of simulations for a buffered system with a buffer size equal to 0.5 Mb, i.e., the total average burst generated by all background sources when they are active. The observed behavior is reported in Table III,[4] in which R_2/R_1, ρ_1/ρ_2, and both R_2/R_1 and ρ_1/ρ_2 are increased from

[4]Note that the 95% confidence intervals for some simulation runs are still relatively large (in the order of 40%), which explains some of the "irregularities" seen in the table.
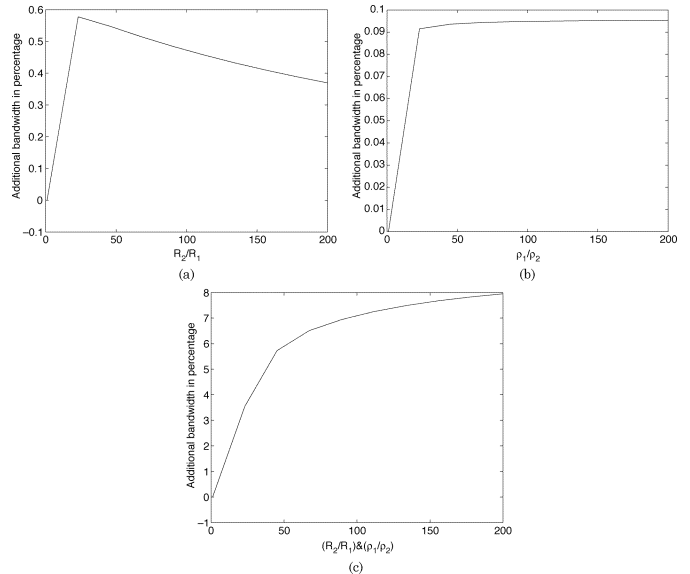

Fig. 4. Additional bandwidth for amny ON-OFF sources. (a) Varying R_2. (b) Varying ρ_2. (c) Varying R_2 ρ_2.

TABLE III
MANY ON-OFF SCENARIOS: 0.5-Mb BUFFER

R_2/R_1	20	50	80	110	140	170	200
P_L^2/P_L	1.66	1.87	1.83	1.80	1.76	1.74	1.70
$\frac{\Delta C}{C}$ (%)	< 0.5						
ρ_1/ρ_2	20	50	80	110	140	170	200
P_L^2/P_L	1.05	1.10	1.09	1.07	1.10	1.14	1.03
$\frac{\Delta C}{C}$ (%)	< 0.1						
$\frac{R_2}{R_1} = \frac{\rho_1}{\rho_2}$	20	50	80	110	140	170	200
P_L^2/P_L	5.25	19.2	42.4	72.1	107.5	144.4	176.3
$\frac{\Delta C}{C}$ (%)	0.06	0.12	0.25	0.64	2.17	4.97	12.0

20 to 200. From Table III, it can be seen that the deviation behavior in a buffered system is qualitatively similar to that in a bufferless system. In particular, significant differences in either peak rates or utilizations again only induce minor performance deviation; while deviation caused by a combination of substantially different peak rates and utilizations is still large, in spite of a 0.5-Mb buffer. This indicates that when aggregating ON-OFF sources, merely increasing the buffer size helps improve overall performance, but does not necessarily eliminate performance differences across users.

B. Many Periodic Sources

In this subsection, we investigate the potential performance deviations when aggregating many periodic sources. Intuitively, we expect that the "smoother" nature of periodic sources together with the large number of sources, will result in relatively small performance deviations for those cases.

The configuration used in this section consists of 1000 type 1 sources with period D_1 and a single type 2 source with period D_2. The ratio D_2/D_1 is then varied from 40 to 200. The link load is first fixed at 0.7 and the buffer size is set to ensure an overall loss probability of 10^{-4}.

The results of those experiments are reported in Table IV. They confirm our expectation that multiplexing a large number

of relatively smooth periodic sources results in only minor performance deviations. The differences in loss probability are of the order of 2% and the amount of additional bandwidth required to bring the loss probability of the type 2 source on par with its target is at most 0.1%. Intuitively, this is because when the number of sources is large, packet arrival epochs are more likely to be randomly distributed, which minimizes the potential variations of the queue length. This makes the loss probability of the type 2 source mostly insensitive to the frequency with which it samples the queue. More formally, this can be deduced from the well-known PASTA result, as the increasing number of sources result in an overall arrival process that approaches Poisson. In such cases, the queue statistics seen by both types of sources should be close to the exact overall system statistics.

The main significance of this finding is that it lends some validity to the use of aggregate QoS solutions for supporting constant bit rate services. Clearly, there are aspects that the periodic model does not capture, e.g., how interactions between flows affect the periodic nature of the traffic as it traverses the network (see [3], [11], and [16] for relevant investigations of this issue). However, it helps confirm that multiplexing CBR sources, even when they have different rates, is reasonably safe in that all sources should approximately see the same performance, at least when many such sources are multiplexed. Note that as is shown in Section IV-B, this does not necessarily hold when only two sources are multiplexed, as a lower rate (longer period) source will often see much higher losses. Understanding how the deviation behaviors evolve as the number of multiplexed sources increases is the topic of the next section.

VI. INTERMEDIATE CONFIGURATIONS

Because of the differences that were observed between the two-source and the many-source configurations, it is of interest to investigate the transition from one set of behaviors to the other. Note that those differences were anticipated, as the two-source and many-source cases were chosen because we expected them to provide results that would be applicable to "small" and "large" systems, respectively, and thus to possibly exhibit different deviation behaviors. What we wish to undertake in this section, is to gain a better understanding of when the results for "small" and "large" systems are applicable.

The approach we take is to test a number of "intermediate" scenarios. The first such scenario involves ON-OFF sources and is tested through simulations. In particular, the number of background (type 1) sources in the system is increased from 1 to 100, while still keeping a single type 2 source.[5] When there is only one type 1 source, its parameters are: $R_1 = 10^8$ bits/s, $b_1 = 0.005$ s, and $\rho_1 = 0.5$. When the number of type 1 sources increases, b_1 and ρ_1 remain fixed but R_1 is decreased so that the total mean rate of the background traffic is kept constant. For the sake of brevity, we only focus on scenarios where the behaviors of the two-source and many-source scenarios differ substantially, i.e., cases that involve differences in ρ and in both R and ρ, and select a large ratio of 200 between the parameters of the two types of sources. Simulations were conducted for

[5]See [26] for another set of intermediate configurations where both the number of type 1 and type 2 sources are varied, while the total number of sources is fixed.

TABLE IV
IMPACT OF SOURCE PERIOD: THE MANY SOURCE CASE

D_2/D_1	40	80	120	160	200
P_L^2/P_L	1.0235	1.0238	1.0239	1.0240	1.0240
$\Delta C/C\%$	0.0981	0.0997	0.1004	0.1004	0.1004

TABLE V
INTERMEDIATE ON-OFF SCENARIOS

	N_1	1	10	20	50	100
$\frac{\rho_1}{\rho_2} = 200$	P_L^2/P_L	100	3.3	1.8	1.1	1.4
	$\Delta C/C$ (%)	23.9	2.5	0.9	0.1	0.07
	N_1	1	10	20	50	100
$\frac{\rho_1}{\rho_2} = 200$	P_L^2/P_L	2	10.4	19.5	44	81
$\frac{R_2}{R_1} = 200$	$\Delta C/C$ (%)	0.02	0.8	1.7	4.4	9.1

both bufferless and buffered systems (the buffer size was again 0.5 Mb). Since results were essentially similar, we only report those of the buffered system simulation.

The results are reported in Table V from which we can draw a number of conclusions. First, when the two types of sources differ only in their utilizations, the convergence to the many-source results is rather rapid. For example, a number of only 50 background sources is sufficient for relying on the many source results. The deviation in loss probability is only 10% and allocating an additional 0.1% of bandwidth eliminates it. This means that when sources differ only in their mean rates and not in their peak rates, e.g., because users are connected to the network through the same kind of access links, aggregating even a relatively small number of (ON-OFF) sources ensures reasonably homogeneous performance.

The situation is somewhat different when dealing with sources that differ in both their utilizations and peak rates. In such a case, the transition from the two-source to the many-source case is much more progressive. Nevertheless, even if it takes a large number of sources before observing the very large performance deviations of the many-source regime in Section V-A, significant differences in performance can already be seen with a relatively small number of sources. Specifically, with only 10 (type 1) sources the loss probability ratio is already greater than 10. This implies that aggregating even a small number of sources that differ in both their peak rates and utilizations can result in substantial deviations.

The second intermediate configuration tested corresponds to the $\sum_j (N_j \cdot D_j + D_i)/D/1$ queue of Section V-B. For this experiment, we still rely on (8) and (9) of Section III-D. We vary the number of type 1 sources from 1 to 100, while fixing the ratio D_2/D_1 at 200. The loss probability ratio and the additional bandwidth needed to satisfy the performance requirements of the type 2 source are given in Fig. 5(a) and (b), respectively.

As can be seen from the figure, the behavior of the many-source scenario is reached rather rapidly, and a background traffic of 100 sources is more than sufficient for that purpose.

VII. APPLICATIONS TO REAL TRAFFIC SOURCES

A. Traffic Aggregation Guideline

What we have observed so far demonstrates that the behavior of loss deviation is primarily affected by the number of sources aggregated. In particular, in the many-source case, the main
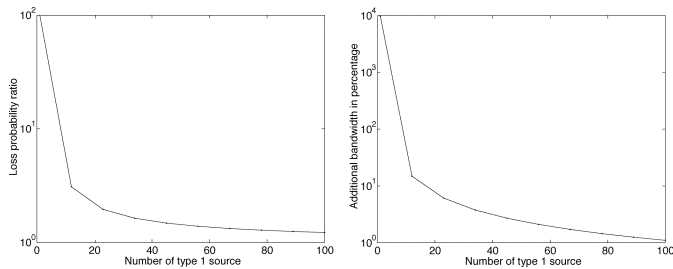
Fig. 5. Intermediate scenarios for periodic sources.

factor responsible for causing significant deviation in performance is a *combination* of high peak rate and low utilization. In contrast, in the two-source case, utilization *alone* can induce significant deviations. What these two cases have in common is that for an individual source to experience much higher losses than the overall loss target, it needs to not only contribute significantly to the onset of congestion when active, but also to do it in a way that does not trigger the allocation of sufficient additional resources. In the case of two (or a few) sources with similar peak rates, each source has a significant influence on congestion whenever it is active. Therefore, decreasing the utilization (lowering its mean rate) of even a single (tagged) source, translates into fewer congestion periods and, better overall performance. This then triggers a decrease in the amount of resources allocated to meet the desired (aggregate) loss probability target. However, this decrease in allocated resources means that the tagged source, whenever it is active, now sees higher levels of congestion and higher losses. Hence, the performance of the tagged source keeps on worsening as its utilization gets smaller. In contrast, when there are many sources with similar peak rates, each source has only a minor effect on congestion when active. In this case, decreasing the utilization of just one tagged source is not sufficient to affect the overall performance, and therefore change the allocated resources. As a result, the performance of the tagged source remains essentially similar to that of the other sources. In order for performance deviations to occur and remain undetected by the system, both the peak rate and the utilization of the tagged source need to be changed. A high peak rate means that the tagged source will typically experience significant congestion when active, while a low utilization indicates that the impact of such congestion on the overall performance is small enough not to trigger the allocation of additional resources. It is this combination of higher peak rate and lower utilization that together results in significant differences in loss probabilities.

We summarize in the following statement:

An individual user will experience a significantly larger loss probability than the overall loss probability of its service class, and thus should not be included in this service class, *if and only if* it satisfies *both* of the following conditions:

1) The traffic contribution of the user, when active, should be capable of inducing substantial congestion.
2) The absolute losses contributed by this user should be small compared to the total traffic, so as not to affect the overall loss probability in a significant way. In other

words, the traffic contribution of the user should be small.

The above statement can be easily cast in the form of a general guideline to determine which sources should and should not be multiplexed. In the next section, based on this guideline, we develop a simple methodology for *quantitatively* identifying when sources can be multiplexed.

B. Methodology for Identifying Dangerous Traffic Mixes

Despite its simplicity and generality, the above traffic aggregation guideline is qualitative in nature and, therefore, difficult to apply directly in practice. This is compounded by the fact that it relies on several assumptions, e.g., Markovian or periodic sources, which often do not hold in practice. In this subsection, we develop a simple methodology that incorporates the above guideline and is able to predict the magnitude of performance deviations for real traffic sources. Based on the predictions of this methodology, decisions can be made on whether or not to multiplex traffic sources. We establish the efficiency of the proposed methodology by testing it for different configurations using the voice and video sources described in Section II-A2.

1) Mapping Real Sources Onto ON-OFF Sources: A first key step of our methodology is to determine how to map real traffic sources onto the analytical source models developed in Section III. Specifically, voice and video sources usually create ON-OFF traffic patterns, which makes a Markov ON-OFF source a natural choice for representing them. There are many possible options for mapping a voice or video source onto an "equivalent" Markovian source, especially if one wants to account for its higher order statistics. Our initial approach, however, is limited to accounting for only first order statistics. Our rationale is that based on the experience derived from the many scenarios we explored using analytical models, we expect first order statistics to play a dominant role in inducing significant performance deviations. This is not to say that higher order statistics have no impact, but only that we expect first order statistics to capture many if not most of the traffic characteristics that can introduce significant performance deviations. As will be shown in the rest of this section, this initial assumption was mostly confirmed by the various experiments performed.

The first traffic source we attempt to map to an ON-OFF source is a raw video source that generates variable size frames at regular time intervals (40 ms for the trace we use). In our study, we assume that a *whole* frame arrives instantaneously, so that the access link serves as a shaper, bounding the maximum rate of the output traffic by its own transmission rate.[6] The shaped traffic emanating from the access link can then be modeled using an equivalent ON-OFF source, with

$$R = R_{\text{link}} \quad b = \frac{F_{\text{avg}}}{R_{\text{link}}} \quad \rho = \frac{b}{40 \cdot 10^{-3}} \quad (12)$$

where R_{link} is the access link speed and F_{avg} is the average frame size of the video trace.

---

[6]We deliberately set access link speed larger than the maximum frame-level bit rate to avoid the situations in which the link cannot transmit one frame in 40 ms and allow consecutive frames to form an extended burst. We found that if this happens, several extremely long bursts may occur and the delay at the access link could be very large.

As with video sources, the traffic of voice sources can be divided into "ON" and "OFF" periods, corresponding to the talk spurt and silence period in human speech. The mapping of voice sources onto ON-OFF sources can be carried out directly using the numbers shown in Table II, i.e., $R = 64$ Kb, $b = 326$ ms, and $\rho = 0.425$. Since voice traffic is reasonably smooth, i.e., during a burst period the voice source generates multiple packets of 80 bytes with an equal spacing of 10 ms, the peak rate of the equivalent ON-OFF voice source is fixed at 64 kb/s instead of equal to the access link speed.[7] As we shall see later, this relative smoothness prevents the voice source from experiencing major performance degradations despite its relatively low rate, even when aggregated with video sources.

2) Decision Methodology: Once video and voice sources have been mapped onto Markov ON-OFF sources, we can apply either (1) and (2) (in the two-source case) or (6) and (7) (in the many-source case) to evaluate the maximum loss probability ratio PLR across all users. The value of the maximum loss probability ratio is then used as a basic measure for deciding whether the corresponding mix of sources is safe or not. Associated with such a decision is a "threshold" region $[\mathrm{PLR}_{\min}, \mathrm{PLR}_{\max}]$, which is used to separate risky traffic mixes, i.e., mixes that give rise to large deviations, from safe ones. Specifically, if the loss probability ratio is less than $\mathrm{PLR}_{\min}$, the traffic mix is considered "definitely safe"; if it is greater than $\mathrm{PLR}_{\max}$, the traffic mix is deemed "definitely dangerous"; if it falls in-between those two values, i.e., inside the threshold region, the traffic mix is flagged as "potentially risky" and requiring additional investigations. We believe that introducing such a "grey" region provides greater flexibility, and can help minimize the impact of inaccuracies inherent in our evaluation methodology.[8] Our hope is that a reasonably narrow threshold region is sufficient to filter out the various traffic mixes for which the model might provide inaccurate answers.

In order to evaluate the efficacy of our methodology in identifying dangerous traffic mixes, we introduce a "consistency test" in which the decision regarding whether or not to aggregate based on the model prediction is compared with the corresponding decision according to simulation. We have conducted extensive consistency tests for many different configurations. Due to space constraints, we only report the most important cases, while complete results can be found in [27]. In all simulations, the packet size of the voice traffic is fixed at 80 bytes and the packet size of the video traffic is fixed at 100 bytes, so that a video frame with a larger size is fragmented into multiple packets. Moreover, the buffer size is fixed at 5000 bytes and the threshold region is always set as [3,5], a reasonably narrow interval. Similar to our previous investigations to the analytical models, we again focus on one of the two types of sources, and vary its parameters so that it experiences a larger loss probability.

Fig. 6.   Loss probability ratio when aggregating two video sources.

C. Testing Our Methodology Using Voice and Video

In this subsection, we test the performance of our methodology using a number of case studies. We again start with the two boundary configurations, i.e, the two-source case and the many-source case.

1) Aggregating Two Sources: The first case we consider in the two-source scenario involves multiplexing a Jurassic low source and a Jurassic high source, which corresponds to two video streams with very different qualities and, therefore, bit rates. We select the lower bit rate source, i.e., the Jurassic low source, as our candidate for experiencing a larger loss probability by varying (increasing) its access link speed, so that its peak rate also increases correspondingly.

The loss probability ratio is shown in Fig. 6. From Fig. 6, we can see that as the access link speed increases, performance deviation can arise with a loss probability ratio that can ultimately reach a value of about 6. More important for us, the figure shows that our methodology accurately tracks both the trend and the magnitude of this deviation. In particular, whenever that actual loss probability ratio is outside the threshold region $[\mathrm{PLR}_{\min}, \mathrm{PLR}_{\max}]$, so is our estimate. Hence, we are able to correctly identify configurations that are either "definitely safe" or "definitely dangerous". Despite the large difference between the utilizations of the two sources, the range of possible loss probability ratios is rather limited (it never exceeds 6). This is because the lower bit rate source, i.e., the Jurassic low source, still contributes a nonnegligible amount of traffic to the overall traffic, which limits the extent to which its performance can deviate from the overall performance. From Table I, a Jurassic low sequence has a mean rate equal to 153.7 Kb, which is about 16.7% of the total input traffic, resulting in a loss probability ratio satisfying $\mathrm{PLR} \leq 6$ according to (11).

A second two-source scenario we consider involves one voice source and one video source, either Jurassic low or Jurassic high. From Table I and Table II, we can see that the voice source has a very low transmission rate compared to either the Jurassic low or the Jurassic high source. This means, according to (11), a relatively large upper bound for PLR, which could indicate a potentially dangerous traffic mix. However, the actual loss probability

---

[7]The peak rate will be equal to the access link speed only when it is less than 64 kb/s, which is lower than the values assumed in this paper.

[8]Ideally, the threshold region should be as narrow as possible, but depending on the accuracy of the model, too narrow a region may either yield a large number of false alarms or fail to properly identify harmful configurations. A wider threshold region reduces those problems at the cost of some imprecision in terms of classifying certain scenarios.
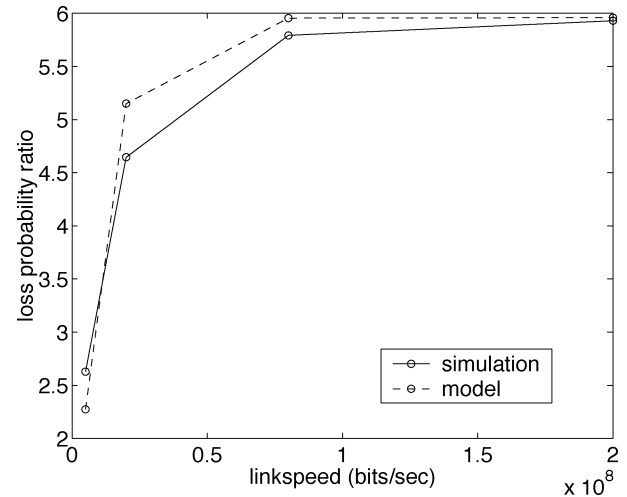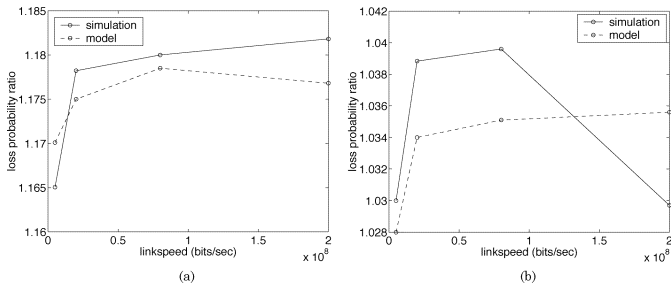
Fig. 7.   Loss probability ratio when aggregating one video source and one voice source. (a) Aggregating voice source and Jurrasic low. (b) Aggregating voice source and Jurrasic high.

ratio is found to be quite small. This is because of the reasonably smooth nature of the voice source, which makes it mostly immune to performance degradations when multiplexed with either of the two video sources. This is consistent with our earlier guideline, since although the voice source satisfies its second condition, i.e., it is a minor traffic contributor, it does not satisfy the first. Hence, it is now the video source that experiences somewhat poorer performance. The magnitude of this performance degradation is, however, very limited because the video source is the major contributor to the overall transmission rate (11) gives an upper bound of 1.04 and 1.18 for the loss probability ratio of Jurassic low and Jurassic high, respectively). As a result, our main concern in these scenarios is whether the model will generate a false alarm, i.e., falsely predict that mixing a voice and a video source is potentially dangerous.

The corresponding results are shown in Fig. 7. From the graphs, we can see that for both Jurassic low and Jurassic high, the model tracks the actual loss probability ratio reasonably accurately. More important, it will not generate any false alarm and will consistently identify the traffic mixes as safe.

2) Aggregating Many Sources: In this section, we explore the case where many voice sources and one video source are aggregated. From the guideline of Section VII-A, we know that for one or more sources to experience substantially worse losses than other sources, its overall traffic contribution should be small while still being able to induce congestion when active. In order to construct such a situation, we choose 500 voice sources as the background traffic and one video source, Jurassic low, as the source of special interest.[9] The loss probability ratio for this configuration is shown in Fig. 8, where we again vary the access link speed of the Jurassic low source so that it experiences a larger loss probability.

As is seen in Fig. 8, the video source's loss probability ratio grows with its access link speed. This evolution is accurately captured by the model even if it slightly overestimates the actual value of the loss probability ratio. Such overestimation, however, does not affect the model's ability to succeed in the consistency test. The figure also shows that in comparison with the two-source scenarios, the range of possible deviations is substantially larger. This is not unexpected, as the traffic contributed by the minor traffic contributor, i.e., the Jurassic low source, to

[9]The choice of Jurassic low is primarily motivated by its lower rate that makes it easier to ensure that its contribution to the overall traffic is relatively small. A similar situation could have been created for the Jurassic high source by increasing the number of background voice sources.
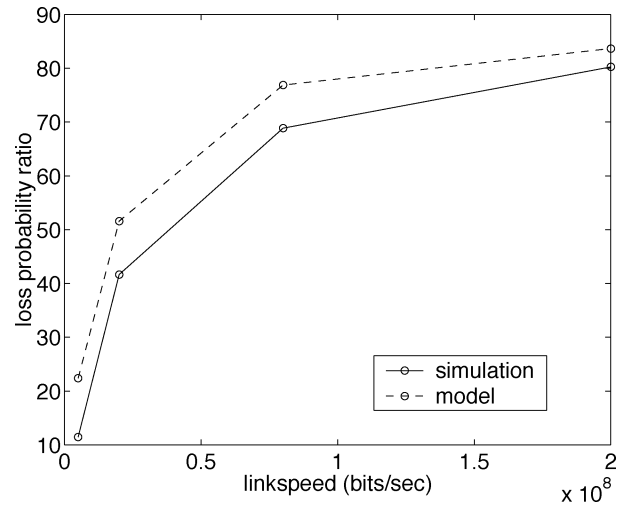

Fig. 8.   Loss probability ratio when aggregating 500 voice sources and one video sources (Jurassic low).
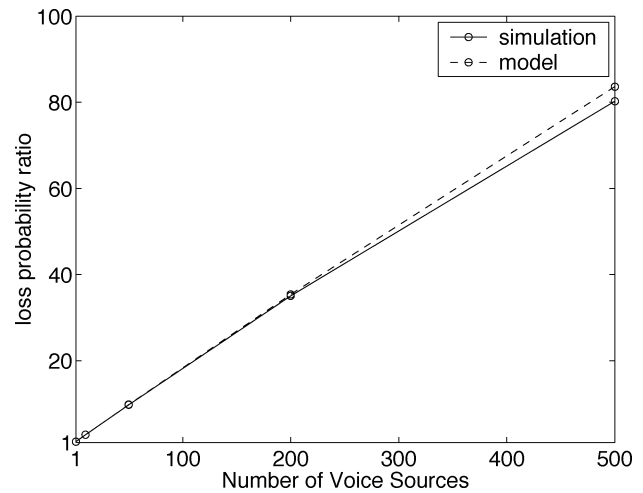

Fig. 9.   Aggregating video sources and Jurassic (low): increasing number of voice sources.

the total traffic volume is much smaller (only 1.12%) than that in the two-source scenarios. This very small contribution to the overall traffic translates into an upper bound of 89.5 for the loss probability ratio, and we can see from the graph that this value is almost achieved when the access link speed is high.

3) Intermediate Configurations: After exploring the behaviors of loss deviations in the two boundary configurations, we again proceed to investigate how one transitions from one set of behaviors to the other as the number of traffic sources varies. Specifically, the scenario we study consists of a variable number (from 1 to 500) of voice sources with one video source (Jurassic low). This allows us to explore the evolution from a safe, two-source traffic mix to an unsafe, many-source mix. The corresponding loss probability ratio is plotted in Fig. 9.

As can be seen from the figure, the model is quite accurate in predicting the loss probability ratio across all scenarios. In addition, the figure shows that the transition from safe to unsafe regions is progressive, roughly following a linear function, but with a relatively steep slope. In particular, mixing the video source with just 50 voice sources already yields a loss

probability ratio close to 10. This is consistent with what we observed in the intermediate configurations involving Markov ON-OFF sources in Section VI, where we had already seen the relatively fast transition from the two-source case to the many-source case. This highlights the fact that it is better to err on the side of caution when considering mixing traffic sources with rather disparate peak and mean rates.

D. Discussions

The data reported in this section only covers a small portion of all the tests that were conducted. More extensive results can be found in [27]. In [27], we observed that in a few cases, especially cases involving aggregating a small number of sources with very different mean rates, our prediction of the loss probability ratio could be off by a nonnegligible amount. However, even in those cases our methodology was typically still able to qualitatively capture the trend in performance deviations and more important, accurately identify dangerous traffic mixes. In other words, while we were off in predicting the absolute value of the loss probability ratio, we were still in the correct safe or dangerous "region." Moreover, the reliability and robustness of the proposed methodology was further established in a more quantitative manner through a worst case analysis. It was shown particularly that the worst case error in predicting the loss probability ratio was small across a broad range of source characteristics, *independent* of the errors introduced when mapping a real traffic source onto an equivalent Markov ON-OFF source. We refer again to [27] for a full presentation of this worst case analysis.

VIII. CONCLUSION

This paper is concerned with an environment where QoS is provided through coarse mechanisms such as service classes, and where guarantees are provided at the class level and enforced through provisioning. Our goals were to determine if and when aggregate performance, in particular loss probabilities, is a good predictor of individual performance, and to identify scenarios where this may not hold. In that context, this paper's contributions are threefold.

First, we derived a number of analytical models that allow the evaluation of individual loss probabilities in environments that only provide aggregate guarantees. These models cover both Markov ON-OFF and periodical sources and are developed for two different boundary configurations, i.e., the two-source scenario and the many-source scenario, so that they are expected to help characterize the possible range of the deviation behavior. Though built on previous works, these models do represent new contributions. In particular, the results of Sections III-A and III-B provide some simple yet useful insights into the impact of individual traffic parameters in the two-source case. Section III-D also provides new results.

Second, we identified a number of configurations that can introduce significant deviations across individual and aggregate loss probabilities, and use these results to derive recommendations and guidelines for avoiding such situations. In particular, we showed that when service classes are used to aggregate only a small number of users, the rate of individual sources

is a key factor and one should avoid multiplexing users with significantly different rates. This holds for both the ON-OFF and the periodic models, although this effect rapidly decreases as more users are being multiplexed. For example, we found that even a difference in rate of 200 results in no more than a factor of 2 in difference in loss probabilities, when 20 CBR sources are multiplexed. When the number of sources being aggregated is large, one should avoid multiplexing flows that differ greatly in both their peak and mean rates, at least in environments where provisioning has some built-in assumptions regarding multiplexing efficiency, i.e., environments where the ON-OFF model is relevant. If a more conservative, e.g., peak rate based, provisioning strategy is used, then multiplexing flows with different traffic characteristics (periods) is reasonably safe, as long as the number of flows aggregated in the service class is sufficiently large. Furthermore, the differences identified between the provisioning environments corresponding to the ON-OFF and periodic models, also carry over to how fast the many-source model becomes applicable. In particular, in the periodic model the impact of rate differences disappears rapidly as the number of sources increases. However, the same does not hold for cases where ON-OFF sources have rather different peak and mean rates. In those cases, a significantly larger number of sources is needed before the characteristics of the many-source model become prominent.

Finally, we developed a simple methodology based on our analytical results, for determining whether aggregation is safe or not when dealing with real-world traffic sources. We demonstrated the efficacy and robustness of the proposed methodology through a series of case studies involving the aggregation of a number of voice and video sources.

In summary, we believe that the results reported in this paper make a valuable contribution to our understanding of the capabilities and limitations of aggregate QoS solutions. In particular, the theoretical framework we have developed can help gain a better insight into why and when performance deviations arise. In addition, the methodology we developed is simple yet accurate enough to provide a practical tool for identifying potentially dangerous traffic mixes.

REFERENCES

[1] MPEG-4 and H.263 Video Trace Library [Online]. Available: http://www-tkn.ee.tuberlin.de/research/trace/trace.html

[2] D. Anick, D. Mitra, and M. M. Sondhi, "Stochastic theory of a data-handling system with multiple sources," *Bell Syst. Tech. J.*, vol. 61, no. 8, pp. 1871–1894, Oct. 1982.

[3] J. C. R. Bennett, K. Benson, A. Charny, W. F. Courtney, and J.-Y. Le Boudec, "Delay jitter bounds and packet scale rate guarantee for expedited forwarding," presented at the IEEE INFOCOM, Anchorage, AK, Apr. 2001.

[4] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, "An Architecture for Differentiated Services," IETF, Request For Comments (Proposed Standard) RFC 2475, Dec. 1998.

[5] A. Charny, J. C. R. Bennet, K. Benson, J.-Y. Le Boudec, W. Courtney, S. Davari, V. Firoiu, and D. Stiliadis, "An Expedited Forwarding PHB(Per-Hop Behavior)," IETF, Request For Comments (Proposed Standard) RFC 3246, (Obsoletes RFC 2598), B. Davie, Ed., Mar. 2002.

[6] H. F. P. Fitzek and M. Reisslein, "MPEG-4 and H.263 video traces for network performance evaluation," *IEEE Network Mag.*, vol. 15, no. 6, pp. 40–54, Nov. 2001.

[7] R. J. Gibbens and P. J. Hunt, "Effective bandwidths for the multitype UAS channel," *QUESTA*, vol. 9, no. 1–2, pp. 17–27, Oct. 1991.

[8] R. J. Gibbens, F. P. Kelly, and P. B. Key, "A decision-theoretic approach to call admission control in ATM networks," *IEEE J. Select. Areas Commun.*, vol. 13, no. 6, pp. 1101–1113, Aug. 1995.

[9] R. Guérin, H. Ahmadi, and M. Naghshineh, "Equivalent capacity and its application to bandwidth allocation in high-speed networks," *IEEE J. Select. Areas Commun.*, vol. 9, no. 7, pp. 968–981, Sep. 1991.

[10] R. Guérin and V. Peris, "Quality-of-service in packet networks: basic mechanisms and directions," *Comput. Netw.*, vol. 31, no. 3, pp. 169–179, Feb. 1999.

[11] R. Guérin and V. Pla, "Aggregation and conformance in differentiated service networks: a case study," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 31, no. 1, pp. 21–32, Jan. 2001.

[12] L. Gün and R. Guérin, "Bandwidth management and congestion control framework of the broad-band network architecture," *Comput. Netw. ISDN Syst.*, vol. 26, no. 1, pp. 61–78, Sep. 1993.

[13] S. Jamin, P. Danzig, S. Shenker, and L. Zhang, "A measurement-based admission control algorithm for integrated services packet networks," *IEEE/ACM Trans. Netw.*, vol. 5, no. 1, pp. 56–70, Feb. 1997.

[14] W. Jiang and H. Schulzrinne, "Analysis of on-off patterns in VoIP and their effect on voice traffic aggregation," presented at the 9th IEEE Int. Conf. Computer Communication Networks, Las Vegas, NV, Oct. 2000.

[15] V. Kumar, T. V. Lakshman, and D. Stiliadis, "Beyond best-effort: router architectures for the differentiated services of tomorrow's Internet," *IEEE. Commun. Mag.*, vol. 36, no. 5, pp. 135–150, May 1998.

[16] L. Massoulié, "Large deviations orderings of point processes in some queueing networks," *Queueing Syst.*, vol. 28, no. 4, pp. 317–335, 1998.

[17] D. Mitra, "Stochastic theory of a fluid model of producers and consumers coupled by a buffer," *Adv. Appl. Prob.*, vol. 20, pp. 646–676, 1988.

[18] I. Norros and J. Virtamo, "Who loses cells in the case of burst scale congestion," in *Proc. 13th Int. Teletraffic Congress (ITC-13)*, Copenhagen, Denmark, Jun. 1991, pp. 59–64.

[19] J. W. Roberts, U. Mocci, and J. Virtamo, Eds., *Broad-band Network Teletraffic—Final Report of Action COST 242*. New York: Springer-Verlag, 1996, vol. 1155.

[20] J. W. Roberts and J. Virtamo, "The superposition of periodic cell arrival streams in an ATM multiplexer," *IEEE Trans. Commun.*, vol. 39, no. 2, pp. 298–303, Feb. 1991.

[21] H. Schulzrinne, "Voice Communication Across the Internet: A Network Voice Terminal," Dept. Comput. Sci., Univ. of Massachusetts, Amherst, MA, Tech. Rep. 92-50, 1992.

[22] T. E. Stern and A. I. Elwalid, "Analysis of separable Markov-modulated rate models for information-handling systems," *Adv. Appl. Prob.*, vol. 23, pp. 105–139, 1991.

[23] I. Stoica, S. Shenker, and H. Zhang, "Core-stateless fair queueing: a scalable architecture to approximate fair bandwidth allocations in high-speed networks," in *Proc. ACM SIGCOMM*, Vancouver, BC, Canada, Sep. 1998, pp. 118–130.

[24] I. Stoica and H. Zhang, "Providing guaranteed service without per flow management," in *Proc. ACM SIGCOMM*, Boston, MA, Sep. 1999, pp. 81–94.

[25] J. T. Virtamo and J. Roberts, "Evaluating buffer requirements in an ATM multiplexer," in *Proc. IEEE GLOBECOM*, Dallas, TX, 1989, pp. 1473–1477.

[26] Y. Xu and R. Guérin. (2001, Jul.) Individual Versus Aggregate QoS: A Loss Performance Study. Univ. of Pennsylvania, Philadelphia, PA. [Online]. Available: http://einstein.seas.upenn.edu/mnlab/publications

[27] Y. Xu and R. Guérin. (2002, Sep.) On Evaluating Loss Performance Deviation: A Simple Tool and Its Practical Implications. Univ. of Pennsylvania, Philadelphia, PA. [Online]. Available: http://einstein.seas.upenn.edu/mnlab/publications

[28] Z.-L. Zhang, Z. Duan, L. Gao, and Y. T. Hou, "Decoupling QoS control from core routers: a novel bandwidth broker architecture for scalable support of guaranteed services," in *Proc. ACM SIGCOMM*, Stockholm, Sweden, Aug. 2000, pp. 71–83.

Ying Xu (S'00) received the B.E. degree from Tsinghua University, Beijing, China, in 1999, and the M.S. degree in electrical engineering from the University of Pennsylvania, Philadelphia, in 2001. He is currently working toward the Ph.D. degree in the Department of Electrical and Systems Engineering, University of Pennsylvania.

His research interests are high-speed networks, network security, overlay and peer-to-peer networks.

Roch Guérin (S'85–M'86–SM'91–F'01) received the Engineer degree from ENST, Paris, France, in 1983, and the M.S. and Ph.D. degrees in electrical engineering from the California Institute of Technology, Pasadena, in 1984 and 1986, respectively.

He joined the Electrical and System Engineering Department, University of Pennsylvania, Philadelphia, in 1998, where he is the Alfred Fitler Moore Professor of Telecommunications Networks. Before joining the University of Pennsylvania, he spent over 12 years at the IBM T. J. Watson Research Center in a variety of technical and management positions. From 2001 to 2004, he was on partial leave from the University of Pennsylvania, starting a company, Ipsum Networks, that pioneered the concept of protocol participation in managing IP networks. He is on the Technical Advisory Board of France Telecom and Samsung Electronics, and has consulted for numerous companies in the networking area. His research has been in the general area of networking, with a recent focus on developing routing and traffic engineering solutions that are both lightweight and robust across a broad range of operating conditions.

Dr. Guérin served as the editor of the ACM SIGCOMM technical newsletter, CCR, from 1998 to 2001. He chaired the IEEE Technical Committee on Computer Communications from 1997 to 1999, and was an editor for the *Journal of Computer Networks*, the *IEEE Communications Surveys*, the IEEE/ACM TRANSACTIONS ON NETWORKING, the IEEE TRANSACTIONS ON COMMUNICATIONS, and the *IEEE Communications Magazine*, and a guest editor of a IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS issue on Internet QoS published in Dec. 2000. He served as member-at-large of the Board of Governors of the IEEE Communications Society from 2000 to 2002, as General Chair of the IEEE INFOCOM'98 conference, and as Technical Program co-chair of the ACM SIGCOMM2001 conference. In 1994 he received an IBM Outstanding Innovation Award for his work on traffic management in the Broad-Band Services Network Architecture. He has been a member of the Association for Computing Machinery (ACM) since 1998.