On Generalized Processor Sharing with Regulated Multimedia Traffic

Chaiwat Oottamakorn[†], Shiwen Mao[‡], and Shivendra S. Panwar[*]

[†]Institute of Engineering and Resources, Walailak University, Nakhon Si Thammarat 80000, Thailand
[‡]The Bradley Department of Electrical and Computer Engineering, Virginia Tech, Blacksburg, VA 24061, USA
[*]Department of Electrical and Computer Engineering, Polytechnic University, Brooklyn, NY 11201, USA
Email: ochaiwat@wu.ac.th, smao@ieee.org, panwar@catt.poly.edu

***Abstract*— Multimedia traffic is becoming an increasing portion of today's Internet traffic due to the flourish of multimedia applications such as music/video streaming, video teleconferencing, IP telephony, and distance learning. In this paper, we study the problem of supporting multimedia traffic using a Generalized Processor Sharing (GPS) server. We derive tight performance bounds on the backlog and delay for regulated multimedia traffic classes in a GPS system. Our approach is quite general since we do not assume any arriving traffic model or any specific traffic regulator, other than that each traffic flow is deterministically regulated. Such deterministic regulators, as well as approximations of the GPS server, are amenable to implementation and are widely implemented in commercial routers. In addition, our analysis is very accurate and has a much high utilization of the server capacity, since we exploit the independence among the traffic flows for higher statistical multiplexing gains. Numerical examples and simulation results are presented to demonstrate the accuracy and merits of our approach, which is practical and well suited for supporting multimedia applications in the Internet.**

I. Introduction

Multimedia traffic is becoming an increasing portion of today's Internet traffic due to the flourish of multimedia applications such as music/video streaming, video teleconferencing, IP telephony, and distance learning. Applications that generate such data can have very diverse Quality of Service (QoS) requirements. One major concern in the design, implementation, and operation of the Internet is how to provide QoS guarantees for applications with diverse QoS requirements, while achieving high utilization of network resources.

QoS guarantees can be provisioned in the Internet using the architectures described in [1] or [2]. However, due to the advances in Dense Wavelength Division Multiplexing (DWDM) technology, over-provisioning in the network core has become a general practice for many service providers. Nevertheless, we argue that over-provisioning does not necessarily solve the QoS provisioning problem. This is because over-provisioning may not be applicable to all segments of the network, due to technical, regulatory, or capital investment limitations. This makes it difficult to guarantee over-provisioning on an end-to-end basis to meet QoS requirements. In order to guarantee end-to-end performance, QoS mechanisms are still needed for the relatively resource constrained access networks (e.g., wireless access networks), while it may be possible to apply over-provisioning in the core.

Various QoS mechanisms have been developed over the years, such as traffic shaping, admission control, QoS signaling and resource reservation, QoS routing, active queue management, and packet scheduling (see [3] for a comprehensive survey). Leaky bucket-based traffic regulation and Generalized Processor Sharing (GPS) are among the most successful QoS mechanisms, since both of them are not only underpinned by rigorous theoretical analysis [4]–[11], but also widely implemented in commercial routers [12].

GPS is a work-conserving scheduling discipline in which multiple traffic classes share a deterministic server. With GPS, each class is associated with a *weight* and is guaranteed a minimum service rate in proportion to its weight whenever it is backlogged. Furthermore, the residual service of the non-backlogged classes are distributed to the backlogged classes in proportion to their weights. Therefore, GPS is efficient in utilizing and sharing the server capacity (since it is work-conserving and the bandwidth is shared by all classes), while being capable of isolating the classes (since each class is guaranteed a minimum rate, it won't be affected by a misbehaving class). By assigning different weights to the classes, service differentiation can be easily achieved.

GPS has been widely studied under various traffic characterizations and under deterministic settings [6] or stochastic settings [7]–[11]. Generally, the bounds obtained by deterministic GPS analysis are very conservative, since they are derived for the worst case scenario that only occurs with a very low probability [6]. Such *hard* QoS guarantees are unnecessary for many multimedia applications, where a certain level of QoS violation is generally acceptable. On the other hand, although the existing bounds obtained by statistical GPS analysis can achieve a much higher resource utilization than deterministic bounds, the traffic characterizations used in such analysis are usually hard to measure and enforce [7]–[11].

In this paper, we investigate the behavior of a high-speed GPS server under *deterministically regulated* multimedia traffic flows, but in a *stochastic setting*. We present a practical framework for supporting multimedia traffic using GPS servers. Within this framework, multimedia traffic flows are regulated with deterministic regulators (such as the popular leaky bucket regulator or other piece-wise linear regulators), thus greatly simplifying user traffic regulation, monitoring, and enforcement, as in a deterministic GPS analysis [6].

0-7803-8938-7/05/$20.00 (C) 2005 IEEE

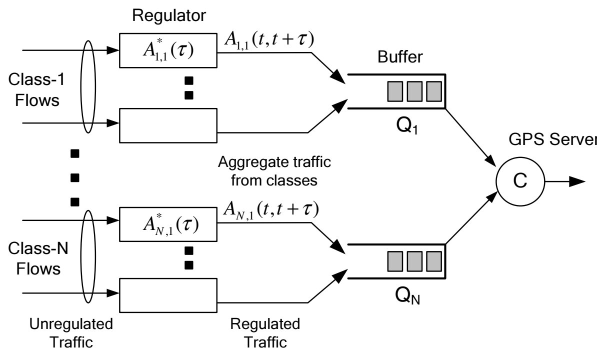

Fig. 1. A GPS System with N classes of regulated multimedia traffic flows.

We then derive tight statistical upper bounds on the tail distribution of buffer occupancy and delay for each class in the GPS system, which achieves a high resource utilization as in a statistical GPS analysis [7]. More specifically, we exploit the independence among the flows and derive tight moment bounds for the aggregate arrival process of each class. Such moment bounds produce an accurate statistical characterization of the aggregate traffic of each class when used with the *Chernoff bound*. Then, we use the notion of feasible ordering as in [6] and [7] to examine the sample path behavior of the multimedia sessions served by a single GPS server in isolation, and derive the backlog (i.e., buffer occupancy) and delay bounds for each class in the GPS system. We also present numerical examples and simulation results to demonstrate the efficacy of the proposed approach.

In summary, our approach has many advantages. First, it obtains tight statistical bounds on the backlog and delay without any assumption on the arriving traffic models. Second, our approach is amenable to implementation and policing as in a deterministic GPS analysis [6], and being capable of achieving high resource utilizations as in a statistical GPS analysis while guaranteeing QoS requirements. Third, our approach is computationally efficient. GPS systems with a large number of classes and hundreds or thousands of flows per class can be easily handled (see Section IV). Such an approach is quite practical and is effective for supporting multimedia applications in the Internet.

The remainder of the present paper is organized as follows. In Section II, we present some preliminary results for the analytical framework. We derive the backlog and delay bounds for a GPS queue with regulated multimedia traffic in Section III. Extensive numerical studies and simulation results are presented in Section IV to demonstrate the accuracy and merits of the proposed framework. Related work is discussed in Section V. We present conclusions and discuss our future work in Section VI.

II. Preliminary Results

We consider a GPS server with a transmission rate C. The server serves N classes of incoming traffic, where each Class i contains a set of flows (denoted as $\mathcal{C}_i$, $i = 1, 2, \cdots, N$). Each flows may itself be an aggregate of the traffic from multiple sessions. As shown in Fig. 1, arriving traffic from each flow k belonging to Class i is regulated by a deterministic envelope $A^*_{i,k}(\tau)$. Arriving class i traffic is buffered at the corresponding queue, i.e., Queue i. In the following we consider a fluid-session traffic model. Note that, without any restriction, our discussion can also be applied to packetized or discrete-time arriving traffic models as well.

A. Regulated Arriving Traffic from a Class

The traffic arriving from a flow k of Class i in an interval (t_1, t_2) is denoted by $A_{i,k}(t_1, t_2)$. We assume that $A_{i,k}(t_1, t_2)$ has the following properties:

(P1) **Additivity:** For any $t_1 < t_2 < t_3$, we have $A_{i,k}(t_1, t_2) + A_{i,k}(t_2, t_3) = A_{i,k}(t_1, t_3)$.

(P2) **Subadditive Bounds:** Traffic is regulated by a deterministic subadditive envelope $A^*_{i,k}$,i.e.,

$$A_{i,k}(t, t+\tau) \le A^*_{i,k}(\tau), \ \ \forall t \ge 0, \forall \tau \ge 0, \quad (1)$$

and

$$A^*_{i,k}(\tau_1) + A^*_{i,k}(\tau_2) \ge A^*_{i,k}(\tau_1 + \tau_2), \ \ \forall \tau_1, \tau_2 \ge 0. \ (2)$$

(P3) **Stationarity:** The A_i's are *stationary* random variables, i.e., $\forall t, \tau, t' > 0$.

$$Pr[A_{i,k}(t, t+\tau) \le x] = Pr[A_{i,k}(t', t'+\tau) \le x], \ \forall x. \quad (3)$$

In other words, the statistical properties of A_i do not change with time.

(P4) **Independence:** All flows are stochastically independent.

Remarks:

- We want to point out that the above assumptions are quite general. The class of subadditive deterministic traffic envelopes is most commonly used as traffic regulators. The assumptions on the randomness of sessions are also quite common. Note that ergodicity is not assumed.
- The traffic regulators most commonly used in practice are *leaky buckets* with a peak rate enforcer. As a generalization of the peak-rate enforced leaky bucket, the conformal traffic on a flow may be characterized by a set of M parameter pairs $\{\sigma^m_{i,k}, \rho^m_{i,k}\}_{m=1,\cdots,M}$, and a deterministic envelope as

$$A^*_{i,k}(\tau) = \min_{m \in \{1,\cdots,M\}} \{\sigma^m_{i,k} + \rho^m_{i,k}\tau\}, \ \ \forall \tau \ge 0, \quad (4)$$

  where the case of $M = 1$ corresponds to the single leaky bucket regulator.
- A consequence of subadditivity of A^*_i is that, for all t,

$$\lim_{\tau \to \infty} \frac{A_{i,k}(t, t+\tau)}{\tau} = \rho_{i,k}. \quad (5)$$

  Eqn. (5) provides an upper bound on the long-term arrival rate for $A_{i,k}$.

B. Generalized Processor Sharing

The GPS scheduler at a network node determines the rates at which traffic in the buffers for the classes is transmitted (see Fig. 1). If there is a backlog at the node, the traffic will be served continuously with a constant rate C, since GPS is a work-conserving scheduling discipline under the assumption that traffic arrivals are characterized by fluid models. As shown in Fig.1, a GPS serves aggregate traffic from N classes. From the definition of the GPS, each Class i is assigned a fixed and real-valued positive number ϕ_i, where $\{\phi_1, \phi_2, ..., \phi_N\}$ is called a GPS weight. The GPS server serves the N classes as follows [6]:

1) Let $S_i(t, t+\tau)$ denote the amount of Class i aggregate traffic served in the interval $[t, t+\tau)$. For a Class i that is backlogged during $[t, t+\tau)$,[1] we have

$$\frac{S_i(t,t+\tau)}{S_j(t,t+\tau)} \geq \frac{\phi_i}{\phi_j}, \qquad j = 1, 2, \cdots, N. \tag{6}$$

2) From Eqn. (6), when any Class i is backlogged, a GPS server guarantees a minimum service rate g_i of [6]

$$g_i = \frac{\phi_i}{\sum_{j=1}^{N} \phi_j} C. \tag{7}$$

3) From the GPS definition, it has been shown that there exists an order among the classes such that, after relabeling the classes [6], [7],

$$\rho_i < \frac{\phi_i}{\sum_{j=i}^{N} \phi_j} \left(C - \sum_{j=1}^{i-1} \rho_j \right), \qquad 1 \leq i \leq N, \tag{8}$$

where $\rho_i = \sum_{k=1}^{\mathcal{C}_i} \rho_{i,k}$, $\rho_{i,k}$ is defined in Eqn. (5), and $\sum_{j=1}^{i-1} \rho_j$ is assumed to be zero when $i = 1$. Note that we assume $\sum_{i=1}^{N} \rho_i < C$ for the system to be stable.

III. Statistical Analysis of a GPS Server with Regulated Traffic Flows

In this section, we will derive the upper bounds on the tail distributions of backlog and delay for each traffic class whose properties for each flow are described in (P1)-(P4). Recall that arriving flows are all independent. In this section, we will proceed with the following two-step approach:

1) We first derive the bound for the moment generating function of the aggregate traffic for each class.
2) Then, we obtain upper bounds on the tail distribution of backlog and delay for each class by using the Chernoff bound and the moment generating function bounds derived in the first step.

The notation used in the following analysis is given in Table I.

TABLE I
NOTATION

$A_i(t, t+\tau)$	aggregate arriving traffic from Class i flows in the interval $[t, t+\tau)$
$S_i(t, t+\tau)$	amount of Class i traffic served in the interval $[t, t+\tau)$.
$A_{i,k}(t, t+\tau)$	arriving traffic from the kth flow of Class i in the interval $[t, t+\tau)$.
$A^*_{i,k}(\tau)$	envelope process of $A_{i,k}(t, t+\tau)$.
$M_{i,k}(\theta, \tau)$	moment generating function of $A_{i,k}(t, t+\tau)$.
$M_i(\theta, \tau)$	moment generating function of the aggregate traffic $A_i(t, t+\tau)$.
C	capacity of the GPS server.
ϕ_i	GPS weight of Class i.
$\rho_{i,k}$	long term average rate of $A_{i,k}(t, t+\tau)$.
ρ_i	long term average rate of $A_i(t, t+\tau)$.
$Q_i(t)$	Class i backlog at time t.
$D_i(t)$	delay experienced by Class i traffic arriving at time t.

A. Moment Bounds for Arriving Aggregate from the Classes

For class i, $i \in \{1, 2, \cdots, N\}$, we denote $A_i(t_1, t_2)$ as aggregate traffic from class i flows in an interval (t_1, t_2), i.e., $A_i(t_1, t_2) = \sum_{k=1}^{\mathcal{C}_i} A_{i,k}(t_1, t_2)$. The moment generating functions of the distributions of an arriving process $A_{i,k}(\cdot)$ are defined as:

$$M_{i,k}(\theta, \tau) := E\left[e^{\theta A_{i,k}(t, t+\tau)} \right]. \tag{9}$$

With the independence assumption (P4), we can derive the moment generating function of the class i aggregate traffic as:

$$M_i(\theta, \tau) := E\left[e^{\theta A_i(t, t+\tau)} \right] = \prod_{k=1}^{\mathcal{C}_i} M_{i,k}(\theta, \tau). \tag{10}$$

[1] A class is backlogged during an interval if there is always traffic from that class waiting in the queue during this interval.

We then apply Theorem 1 in [13], which presents a bound on the moment generating function of each flow, and have

$$M_{i,k}(\theta, \tau) \leq 1 + \frac{\rho_{i,k}\tau}{A^*_{i,k}(\tau)} \left(e^{\theta A^*_{i,k}(\tau)} - 1 \right). \tag{11}$$

Substituting Eqn. (11) into Eqn. (10), we derive the bound on the moment generating function of the class i aggregate traffic as follows:

$$M_i(\theta, \tau) \leq \prod_{k=1}^{\mathcal{C}_i} \left[1 + \frac{\rho_{i,k}\tau}{A^*_{i,k}(\tau)} \left(e^{\theta A^*_{i,k}(\tau)} - 1 \right) \right]. \tag{12}$$

Remarks:

- $M_i(\cdot)$ is expressed in terms of simple deterministic envelopes of arriving flows, that is, the leaky bucket or multiple leaky bucket functions (see Eqn. 4).
- If the Class i flows are homogeneous, i.e., all flows in Class i has the same envelope, $A^*_{i,1}(\tau) = A^*_{i,2}(\tau) = \cdots = A^*_{i,\mathcal{C}_i}(\tau)$, then $M_i(\cdot)$ has the following simplified form

$$M_i(\theta, \tau) \leq \left[1 + \frac{\rho_{i,1}\tau}{A^*_{i,1}(\tau)} \left(e^{\theta A^*_{i,1}(\tau)} - 1 \right) \right]^{\mathcal{C}_i}. \tag{13}$$

B. Statistical Bounds on Backlog and Delay

In this section, we study the sample path behavior of the aggregate traffic of each class in the GPS system as shown in Fig. 1. We employ the moment generating function bounds derived in the last subsection to determine the upper bounds on the tail distributions of the backlog and delay for each class.

Consider a tagged Class i. Let time 0 indicate the start of a busy period of length B during which Class i is continuously backlogged. During the busy period, traffic will be served and transmitted continuously at a constant rate C. For some interval $[0, \tau) \in B$, since there is backlog and the work conserving property, the server is busy during this interval.

Then we have

$$\sum_{j=1}^{N} S_j(0,\tau) = C\tau,$$

$$\sum_{j=1}^{i-1} S_j(0,\tau) + \sum_{j=i}^{N} \frac{\phi_j}{\phi_i} S_i(0,\tau) \geq C\tau. \quad (14)$$

We can solve Eqn. (14) for $S_i(0,\tau)$ and have

$$S_i(0,\tau) \geq \frac{\phi_i}{\sum_{j=i}^{N} \phi_j} \left(C\tau - \sum_{j=1}^{i-1} S_j(0,\tau) \right). \quad (15)$$

It has been shown that for regulated classes, the maximum backlog and delay are achieved when all the classes become greedy from time 0 [6], [10]. Since the busy period starts at time 0, the amount of serviced traffic of a class is bounded by its arrival during $[0,\tau)$. The following inequality holds for all classes j:

$$S_j(0,\tau) \leq A_j(0,\tau). \quad (16)$$

Using Eqn. (16), Eqn. (15) can be rewritten as follows:

$$S_i(0,\tau) \geq \frac{\phi_i}{\sum_{j=i}^{N} \phi_j} \left(C\tau - \sum_{j=1}^{i-1} A_j(0,\tau) \right). \quad (17)$$

Define $\psi_i = \frac{\phi_i}{\sum_{j=i}^{N} \phi_j}$. As $Q_i(0) = 0$, we have $Q_i(\tau) = A_i(0,\tau) - S_i(0,\tau)$ for $\tau \leq B$. Then, we obtain

$$Q_i(\tau) \leq A_i(0,\tau) - \psi_i \left(C\tau - \sum_{j=1}^{i-1} A_j(0,\tau) \right). \quad (18)$$

Theorem 1: Given a GPS server with a transmission rate C serving N incoming traffic classes $\{A_1, A_2, \cdots, A_N\}$ satisfying (P1)-(P4). Assume that $\sum_{i=1}^{N} \rho_i < C$ and a feasible ordering with respect to $\{\phi_1, \phi_2, \cdots, \phi_N\}$ and $\{\rho_1, \rho_2, \cdots, \rho_N\}$. For any time instance τ in a busy period and any $q_i \geq 0$,

$$Pr[Q_i(\tau) \geq q_i] \leq \inf_{\theta \geq 0} \left\{ e^{-\theta(q_i + \psi_i C\tau)} \prod_{k=1}^{\mathcal{C}_i} \left[1 + \frac{\rho_{i,k}\tau}{A_{i,k}^*(\tau)} \left(e^{\theta A_{i,k}^*(\tau)} - 1 \right) \right] \prod_{j=1}^{i-1} \prod_{k=1}^{\mathcal{C}_j} \left[1 + \frac{\rho_{j,k}\tau}{A_{j,k}^*(\tau)} \left(e^{\psi \theta A_{j,k}^*(\tau)} - 1 \right) \right] \right\}. \quad (19)$$

Proof: We will use the moment generating function bound in Eqn. (12) to construct the upper bounds on the tail distribution of backlog of Class i. To obtain a rigorous upper bound on $Pr[Q_i(\tau) \geq q_i]$, recall the Chernoff bound for a random variable Y:

$$Pr[Y \geq y] \leq e^{-\theta y} E\left[e^{\theta Y} \right], \ \forall \theta \geq 0. \quad (20)$$

In particular, for $Q_i(\cdot)$ in Eqn. (18), this gives

$$\begin{aligned} & Pr[Q_i(\tau) \geq q_i] \\ \leq\ & Pr\left[A_i(0,\tau) + \psi_i \sum_{j=1}^{i-1} A_j(0,\tau) \geq q_i + \psi_i C\tau \right] \quad (21) \\ \leq\ & e^{-\theta(q_i + \psi_i C\tau)} M_i(\theta,\tau) \prod_{j=1}^{i-1} M_j(\psi_i \theta, \tau), \quad (22) \end{aligned}$$

where $M_i(\cdot)$ is given in Eqs. (12) and (13). Note that Eqn. (22) holds true for any positive value of θ. In order to get a tight bound on the tail distribution, we need to choose a θ that minimizes the right hand side (RHS) of Eqn. (22), i.e.,

$$\begin{aligned} & Pr[Q_i(\tau) \geq q_i] \\ \leq\ & \inf_{\theta \geq 0} \left\{ e^{-\theta(q_i + \psi_i C\tau)} M_i(\theta,\tau) \prod_{j=1}^{i-1} M_j(\psi_i \theta, \tau) \right\}. \quad (23) \end{aligned}$$

■

Next, we derive the upper bounds on the delay tail distribution for each class. Again, consider a busy period of the GPS queue starting at time 0 during which Class i is continuously backlogged. Based on the definition of service curves, it has been shown that for traffic enqueued at any time instance τ within the busy period, the corresponding queueing delay is upper bounded by [14]:

$$D_i(\tau) = \min\{z : z \geq 0 \text{ and } A_i(0,\tau) \leq S_i(0,\tau + z)\}. \quad (24)$$

Eqn. (24) applies to the class of First-Come-First-Serve (FCFS) systems, to which GPS queues belong.

For a given delay bound d_i for Class i, delay violation occurs when there exists any traffic experiencing a delay longer than the delay bound, that is,

$$D_i(\tau) = \min\{z : z \geq 0 \text{ and } A_i(0,\tau) \leq S_i(0,\tau + z)\} \geq d_i. \quad (25)$$

or,

$$A_i(0,\tau) \geq S_i(0,\tau + d_i). \quad (26)$$

Theorem 2: Given a GPS server with a transmission rate C serving N incoming traffic classes $\{A_1, A_2, \cdots, A_N\}$ satisfying properties (P1)-(P4). Assume that $\sum_{i=1}^{N} \rho_i < C$ and a feasible ordering with respect to $\{\phi_1, \phi_2, \cdots, \phi_N\}$ and $\{\rho_1, \rho_2, \cdots, \rho_N\}$. For any time instance τ in a busy period and any $d_i \geq 0$,

$$Pr[D_i(\tau) \geq d_i] \leq \inf_{\theta \geq 0} \left\{ e^{-\theta \psi_i C(\tau + d_i)} \prod_{k=1}^{\mathcal{C}_i} \left[1 + \frac{\rho_{i,k}\tau}{A_{i,k}^*(\tau)} \left(e^{\theta A_{i,k}^*(\tau)} - 1 \right) \right] \prod_{j=1}^{i-1} \prod_{k=1}^{\mathcal{C}_j} \left[1 + \frac{\rho_{j,k}(\tau + d_i)}{A_{j,k}^*(\tau + d_i)} \left(e^{\psi \theta A_{j,k}^*(\tau + d_i)} - 1 \right) \right] \right\}. \quad (27)$$

Proof: From Eqn. (26), we have

$$Pr[D_i(\tau) \geq d_i] = Pr[A_i(0,\tau) \geq S_i(0,\tau + d_i)]. \quad (28)$$

Inserting Eqn. (17) into Eqn. (28), we have

$$\begin{aligned} & Pr[D_i(\tau) \geq d_i] \\ \leq \quad & Pr\left[A_i(0,\tau) + \psi_i \sum_{j=1}^{i-1} A_j(0, \tau + d_i) \geq \psi_i C(\tau + d_i)\right], \\ \leq \quad & e^{-\theta\psi_i C(\tau+d_i)} M_i(\theta, \tau) \prod_{j=1}^{i-1} M_j(\psi_i \theta, \tau + d_i). \end{aligned} \tag{29}$$

■

Remarks:

- If Class i's buffer overflow probability should be less or equal to ϵ_i, the corresponding buffer requirement q_i should be set to the smallest number satisfying the inequality
$$Pr[Q_i(\tau) \geq q_i] \leq \epsilon_i, \tag{30}$$
for any time instance τ in a busy period.
- Similarly, if Class i's delay violation probability should be less or equal to ϵ_i, the corresponding delay d_i is the minimum value that satisfies the inequality
$$Pr[D_i(\tau) \geq d_i] \leq \epsilon_i, \tag{31}$$
for any time instance τ in a busy period.
- For admission control test, Theorems 1 can be used to compute the corresponding loss probability for a given q_i. If the computed probability is less than ϵ_i, then the flows are admissible. Admission control test based on delay can be conducted in the similiar fashion using Theorem 2. For classes with both loss and delay requirements, the flows are admissible only when both requirements are satisfied.
- The RHS of Eqn. 19 and the RHS of Eqn. 27 are both functions of θ. In pactice, the miminum can be computed using standard non-linear optimization techniques, such as quadratic programming.

IV. Numerical Evaluation

In this section, we present a set of numerical experiments to demonstrate the tightness of the derived bounds on the tail distributions of the backlog and delay for each traffic class. We also examine the performance of the derived bounds on admission control tests.

In the experiments, we set $N = 4$ as in the specification of IETF's Diffserv architecture [2]. The four types of traffic used in the experiments, as well as the corresponding parameters, are listed in Table II. More specifically, the first class consists of on-off traffic sources. The sojourn time in the states are exponentially distributed with average on time $= 100$ msec and average off time $= 600$ msec. The second class consists of on-off sources with Pareto distributed sojourn times. The mean on and off times are 100 msec and 325 msec, respectively. In the on state, these on-off sources transmit data at their peak rates; in the off state, the sources do not generate any traffic. The third and fourth classes are MPEG-4 compressed video traces from the movie *Jurassic Park* and a soccer game [15]. The first class of sources generate short-range dependent traffic, which is a good model for IP telephony. The last three classes of sources produce long-range dependent and self-similar traffic in aggregate, which can be used to model other multimedia applications such as video streaming and video teleconferencing. All of the sources are assumed to be regulated by multiple leaky buckets [15]. We consider a high-speed GPS server with a capacity C of 0.8–1.0 Gbps.

TABLE II

Traffic Source Parameters

	Avg. Rate (Kbps)	Peak Rate (Kbps)	GPS Assignment ϕ_i
Class 1: on-off (Exp.)	230	1600	280
Class 2: on-off (Pareto)	520	2220	500
Class 3: Video Trace 1	773	2687	800
Class 4: Video Trace 2	1053	3150	1000

A. Experiment 1: Backlog and Delay Bounds

In this experiment, we compute the derived bounds on the tail distributions of backlog and delay as in Eqs. (19) and (27) and compare them with OPNET simulations results. Specifically, we studied a GPS system with 4 classes, each having 300 flows. The source parameters and the GPS assignments for the classes $\{\phi_i\}_{i=1,2,3,4}$ are given in Table II. For simulations, we implemented a GPS system using the OPNET simulator. During the simulations, it was assumed that traffic arrivals from all classes have different start times, which were randomly chosen. The length of each simulation was 200 seconds of simulated time. In order to illustrate the tightness of the derived bounds, we performed a large number of simulations. For each set of parameters, we performed several simulations and computed the average performance results such as the probabilities of buffer overflow $Pr[Q \geq q]$ and delay violation $Pr[D \geq d]$. The average server utilization is defined as $\rho = \sum_{i=1}^{4} \rho_i / C$, where ρ_i is the long-term average rate of Class i aggregate traffic, $i = 1, 2, 3, 4$.

The results of the experiments are plotted in Figs. 2–3. Figure 2 shows the tail distribution of buffer overflow as functions of each class's buffer size q_i and average server utilization ρ, which was varied by changing the server capacity. From Fig. 2, the following can be observed. First, in Fig. 2, our derived bounds on backlog given in Eqn. (19) are tight upper bounds of those obtained from the simulations. For example, in Fig. 2(b), for Class 1 with server utilization $\rho = 85.7\%$ and $q = 50$ Mbits, the derived backlog bound and that from the simulations are 2.84×10^{-4} and 1×10^{-4}, respectively. Second, in Fig. 2, all classes yield better performance as buffer size is increased or average server utilization is decreased (i.e., when C is increased). Finally, in Fig. 2(b), for average server utilization $\rho \geq 88\%$, the higher the class index, the worse the performance. This is because there was not enough server resource (as C is decreased) for the more bursty traffic belonging to Class 2, 3, and 4.

Figures 3 depicts the tail distribution of delay violation as functions of each class's delay bound d_i and average server utilization ρ. Similar observations can be made from these figures. Clearly, better performance is observed as the delay bound is increased or as the average server utilization is decreased. Our derived bounds on the delay are upper bounds on the simulation results and match the simulation results closely, as can be observed from Fig.3.

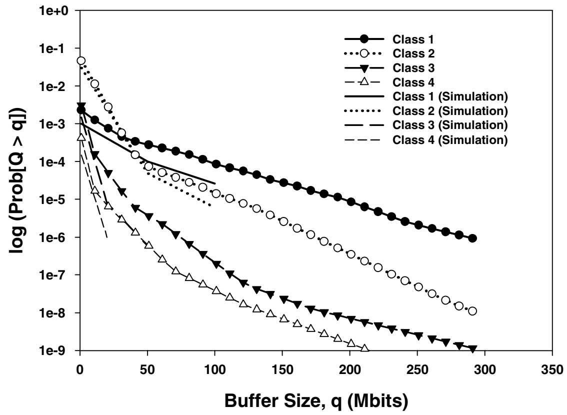

(a) Tail distribution as a function of buffer size q with fixed average server utilization ($\rho = 85.7\%$)

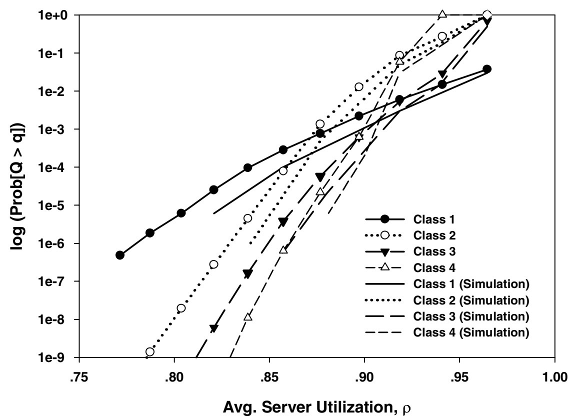

(b) Tail distribution as a function of average server utilization ρ with fixed buffer size ($q = 50$ Mbits).

Fig. 2. Example 1: Tail distribution of buffer overflow for each class as compared with the simulation results.

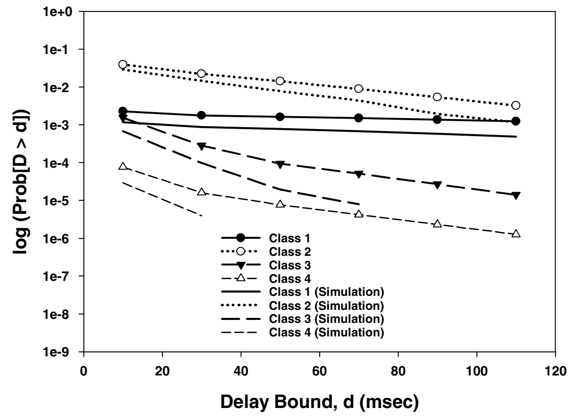

(a) Tail distribution as a function of the delay bound d with fixed average server utilization ($\rho = 85.7\%$)

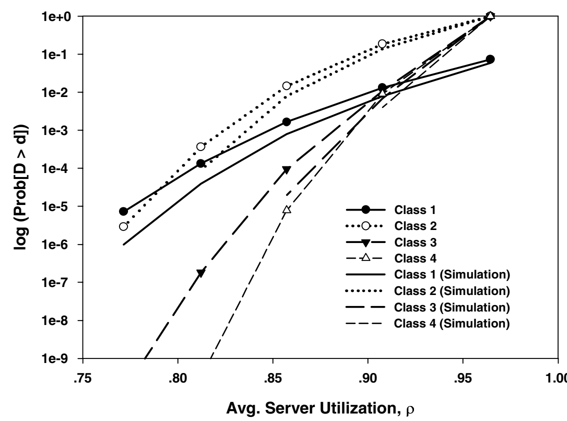

(b) Tail distribution as a function of average server utilization ρ with fixed delay bound ($d = 50$ msec).

Fig. 3. Example 1: Tail distribution of delay for each class as compared with the simulation results.

B. Experiment 2: Admissible Regions

The admissible region of a GPS system is defined as the number of admissible flows of the classes whose buffer overflow and delay requirements are satisfied. By definition, the admissible region indicates the effectiveness of traffic classes in exploiting GPS's resource sharing capability. In each experiment reported in this section, we determine the admissible region for each class based on its traffic characterization and QoS requirements using the bounds in Eqs. (19) and (27).

Consider a GPS server with a capacity of 900 Mbps serving four traffic classes (as listed in Table II) with a mixture of buffer overflow and delay QoS requirements. For an easier presentation, we fixed the number of Class 1 sources at 300, which serves as background traffic. Class 2 required statistical service with $Pr[Q_2 \geq 50Mbits] = 10^{-4}$; Class 3 required statistical service with $Pr[D_3 \geq 30msec] = 10^{-3}$; Class 4 required statistical service with $Pr[D_4 \geq 50msec] = 10^{-4}$.

The resulting admissible region is shown in Fig. 4. We compare the results with average rate allocation.[2] In average rate allocation, for every combination of the numbers of Classes 1, 2, and 3 sources, the number of Class 4 sources is determined by the ratio of the average available capacity over the average rate of a Class 4 source. Figure 4 shows that the admissible region achieved by our bounds is very close to that obtained by average rate allocation, i.e., between $85-89\%$ of the admissible region obtained by average server allocation. This demonstrates the tightness of our backlog and delay bounds, especially when the number of sources is large. Consequently, our approach is able to fully exploit the statistical multiplexing gain of multi-class GPS sharing.

For comparison, we also plot the admissible region obtained by peak rate allocation which is deployed in deterministic analyses. In peak rate allocation, for every combination of the numbers of Class 1, 2, and 3 sources, the number of Class 4 sources is determined by the ratio of the average available capacity over the peak rate of a Class 4 source. As can be seen from Fig. 4, the admissible region using peak rate allocation is a lower bound for that of any QoS provisioning scheme. It is very conservative since statistical multiplexing is not exploited.

[2]Since average rate allocation only guarantees finiteness of delays and average throughput, the admissible region obtained by average rate allocation is an upper bound for that of any QoS provisioning method.

V. Related Work

The exact analysis of a GPS system is non-trivial, since the service rate a class receives is coupled with the backlog

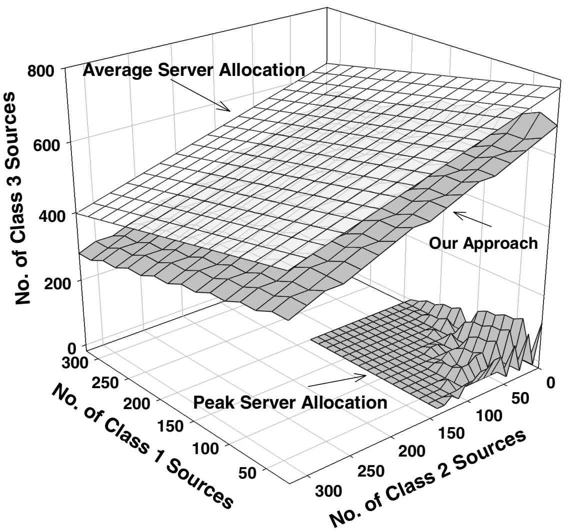

Fig. 4. The admissible region for the three classes in a four-class GPS system.

status and instant rates of all other classes. Over the years, GPS has been widely studied under various traffic characterizations, such as leaky bucket regulated sources [6], exponential bounded burstiness (EBB) sources [7], Markov Modulated Fluid Process (MMFP) sources [11], Gaussian traffic sources [8], and heavy-tailed sources [9], [10]. These papers provide great insights into the behavior of GPS servers under bursty traffic flows. One of the most widely used technique in GPS analysis is the *feasible ordering* or *feasible partitioning* technique. With this technique, a GPS system can be decomposed into a set of separate FIFO queues from which performance bounds can be derived [6], [7]. We also used this technique in the present paper.

As discussed, the bounds obtained by deterministic GPS analysis are very conservative, since worst-case analysis is employed [6], while the existing statistical GPS bounds are not amenable for traffic regulation, monitoring, and enforcement [7]–[11]. In addition, the performance bounds found in [9] are asymptotic bounds that are only valid for very large buffer sizes. Although shedding great insight on the GPS behavior for long-tailed sources, such bounds may not be applicable for multimedia traffic, where the generally tight end-to-end delay requirements prohibit the use of very large buffers in the intermediate routers. Finally, the analysis in [11] has the "state-explosion" problem, making it not suitable for handling a large number of sources, and the bounds get loose as the number of classes increases.

These observations motivated us to investigate the behavior of a GPS server under regulated multimedia traffic flows, but in a stochastic setting. As a result, our approach has the advantage of being as amenable to implementation and policing as deterministic GPS, and is capable of achieving the high resource utilizations achievable under statistical GPS analysis. In addition, our analytical approach is scalable. That is, it can easily handle a large number of classes and achieve tight bounds. It is therefore a practical and efficient approach for supporting regulated multimedia traffic in the Internet.

VI. Conclusions

In this paper, we studied the problem of QoS provisioning for regulated multimedia applications using a GPS server. Based on the general assumptions that the flows are independent and that each flow is deterministically regulated, we derived tight backlog and delay bounds for the GPS system. We also demonstrated the accuracy and merits of our approach using examples and simulation results with MPEG-4 video traces. The framework presented in this paper is quite general and practical for supporting multimedia applications in the Internet.

Acknowledgment

The authors would like to thank the anonymous reviewers for helpful comments. This research has been supported in part by the National Science Foundation under Grants ANI-0081375, and by the New York State Office of Science, Technology and Academic Research (NYSTAR) through the Center for Advanced Technology in Telecommunications (CATT) at Polytechnic University.

References

[1] R. Braden, D. Clark, and S. Shenker, "Integrated services in the internet architecture: an overview," June 1994, IETF RFC 1633.

[2] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, "An architecture for differentiated service," Dec. 1998, IETF RFC 2475.

[3] H. J. Chao and X. Guo, *Quality of Service Control in High-Speed Networks*, John Wiley & Sons, Inc., 2001.

[4] Jean-Yves Le Boudec and Patrick Thiran, *Network Calculus: A Theory of Deterministic Queueing Systems for the Internet*, Springer Verlag, London, Britain, 2004.

[5] C. S. Chang, *Performance guarantees in communication networks*, Springer Verlag, London, Britain, 2000.

[6] A. K. Parekh and R. G. Gallager, "A generalized processor sharing approach to flow control in integrated services networks: the single-node case," *IEEE/ACM Transactions on Networking*, vol. 1, no. 3, pp. 344–357, June 1993.

[7] Z.-L. Zhang, D. Towsley, and J. Kurose, "Statistical analysis of the generalized processor sharing scheduling discipline," *IEEE J. on Selected Areas in Commun.*, vol. 13, no. 6, pp. 1071–1080, Aug. 1995.

[8] P. Mannersalo and I. Norros, "GPS schedulers and gaussian traffic," in *Proc. of INFOCOM 2002*, New York, NY, June 2002, pp. 1660–1667.

[9] S. Borst, O. Boxma, and P. Jelenkovie, "Generalized processor sharing with long-tailed traffic sources," in *Proc. of ITC-16*, Edinburgh, Scotland, June 1999, pp. 345–354.

[10] Flavio M. Pereira, Nelson L. S. Fonseca, and Dalton S. Arantes, "On the performance of generalized processor sharing servers under long-range dependent traffic," Tech. Rep., the Institution of Computer Science, State University of Campinas, IC-01-11, Oct. 2001.

[11] F. Lo Presti, Zhi-Li Zhang, and D. Towsley, "Bounds, approximations and applications for a two-queue gps system," in *Proc. of INFOCOM 1996*, San Francisco, CA, Mar. 1996, pp. 1310–1317.

[12] Cisco Systems, "QC: Cisco IOS Release 12.0 quality of service solutions configuration guide," Cisco IOS Documentation.

[13] R. Boorstyn, A. Burchard, J. Liebeherr, and C. Oottamakorn, "Statistical service assurances for traffic scheduling algorithms," *IEEE Journal on Selected Areas in Communications. Special Issue on Internet QoS*, vol. 18, no. 12, pp. 2651–2664, Dec. 2000.

[14] H. Sariowan, R. Cruz, and G. Polyzos, "Scheduling for quality of service guarantees via service curves," in *Proceedings of the International Conference on Computer Communications and Networks (ICCCN)*, Las Vegas, NV, Sept. 1995, pp. 512–520.

[15] F. Fitzek and M. Reisslein, "MPEG-4 and H.263 traces for network performance evaluation (extended version)," Tech. Rep. Technical Report TKN-00-06, Technical University Berlin, Dept. Electrical Eng., Germany, Oct. 2000.