Bounds, Approximations and Applications for a Two-Queue GPS System

Francesco Lo Presti
Dipartimento di Ingegneria Elettronica
Università di Roma "Tor Vergata"
00133 Roma Italy

Zhi-Li Zhang*, Don Towsley*
Computer Science Department
University of Massachusetts
Amherst, MA 01003, USA

Abstract

In this paper we study the performance of a multiplexer using the *Generalized Processor Sharing* (GPS) scheduling to serve Markov Modulated Fluid Sources (MMFSs). We focus on a two-queue GPS system serving two classes of sources. By using a bounding approach combined with an approximation approach and by taking advantage of the specific structure of MMFSs, we are able to derive a lower bound and an upper bound approximation on queue length distributions for each class of the GPS system. Numerical investigations show that the lower bound and the upper bound approximation are very accurate. Hence our work greatly improves the earlier results on GPS scheduling in [15, 17] which are obtained for a more general stochastic model. Application of our performance bounds to call admission control and bandwidth sharing is also illustrated, and a comparison with FIFO and strict priority in different scenarios is presented. We show that the flexibility provided by GPS does not provide much better performance than FIFO and priority when the classes only have loss requirements. However, this flexibility provides better performance when the classes exhibit delay requirements as well as loss requirements.

1 Introduction

The future broadband ISDN network is expected to provide a variety of services to users, integrating currently separate telephone, cable and data networks. Due to different traffic characteristics and *Quality-of-Service* (QoS) requirements of network traffic, the coexistence of voice, video and data in the same network poses new issues in packet scheduling, call admission control, and bandwidth sharing. For example, data traffic is generally considered to be very bursty and relatively insensitive to delay, but loss-sensitive. On the other hand, real-time traffic such as voice and video is delay-sensitive, but can tolerate some loss. Among real-time traffic, voice is usually less bursty and has a smaller bit rate, whereas video is generally burstier and has a higher bandwidth requirement. This diversity of traffic characteristics and QoS requirements indicates that different classes of traffic should be treated separately according to their respective QoS requirements. Packet scheduling and call admission control are important mechanisms for achieving this separation and, at the same time, allow bandwidth sharing among different classes. The simultaneous need of isolation and bandwidth sharing among classes has been identified in many recent papers (see, *e.g.* [3]). Isolation among classes is important to ensure that misbehavior of traffic in one class will not affect other classes. Bandwidth sharing is not only the result of an integrated service network, but can also be used to exploit statistical multiplexing gain made possible by such a network.

One of the most promising packet scheduling disciplines proposed in recent years is the so-called *Generalized Processor Sharing* (GPS) [10] (also known as Weighted Fair Queueing [4]). One important feature of GPS is its ability to provide isolation among different classes, while, at the same time, allowing bandwidth sharing among these classes. A second feature is that the bandwidth sharing mechanism of GPS is explicitly controllable. Because of these features, GPS is recommended in [3, 12] as a scheduling discipline for networks that must support different service classes.

The performance of GPS has been studied in a variety of settings [10, 11, 15, 17]. Most of these studies are based on the so-called bounding approach where performance bounds on the interested metrics such as loss or delay are derived, usually based on some very general source models, *e.g.* Cruz's *Linear Bounded Arrival Process* (LBAP) traffic model [2] in a deterministic setting, or Yaron and Sidi's Exponentially Bounded Burstiness (E.B.B.) process model [14] in a stochastic setting. Although results derived from the above approaches are applicable to many different arrival processes such as the commonly used Markov Modulated Poisson Processes (MMPPs) and Markov Modulated Fluid Sources (MMFSs), they usually lead to quite loose bounds.

In this paper, we study GPS scheduling with Markov Modulated Fluid Sources. For simplicity, we focus on GPS scheduling with only two service classes, each of which has its own buffer. We call such a system a two-queue GPS system. We combine the bounding approach of [17] with the approximation approach using spectral analysis techniques developed in [1, 5, 6, 7]. We exploit the idea of decomposition first employed in [17]. However, by taking advantage of the specific structure of MMFSs, we are able to derive lower and upper bound systems that better approximate the original GPS system. The decomposition approach leads to the study of single-queue statistical multiplexing systems with modulated service processes.

*The work of these authors was supported in part by the National Science Foundation under grant CCR-9119922.

0743-166X/96 $5.00 © 1996 IEEE

Lower and upper bounds on the queue-length distribution for each queue of the GPS system are easily obtained for these systems. Comparison with simulation shows that our lower and upper bound approximations are generally very close to the simulated values. Refined effective bandwidth approximations based on these bounds are also easily derived. Application of these bounds for GPS scheduling in the context of call admission control and bandwidth sharing is also illustrated, and comparison with FIFO (First-In First-Out) and strict priority in different scenarios is made and discussed.

The rest of the paper is organized as follows. In Section 2, we define GPS scheduling formally and discuss related work on GPS and statistical multiplexing system with MMFSs. In Section 3, we describe the two-queue GPS system model. In Section 4, we derive lower and upper bound approximation for the two-queue GPS system based on study of a statistical multiplexing system with modulated service process. Numerical investigations are presented to show the tightness of the bounds. In Section 5, applications of the bounds to call admission control and bandwidth sharing are illustrated. Finally, the paper is concluded in Section 6.

2 Preliminaries and Related Work

2.1 Generalized Processor Sharing

Generalized Process Sharing (GPS) is a work-conserving scheduling discipline that can be regarded as the limiting form of a weighted round robin policy, where traffic from sessions is treated as infinitely divisible fluid (hence there is no notion of "packet" in this traffic model [10]). Assume we have n sessions sharing a GPS server with rate r. Associated with the sessions is a set of parameters $\{\phi_i\}_{1\leq i\leq n}$ (called the *GPS assignment*) which determines the minimum sharing of bandwidth of each session. Each session is guaranteed a minimum service rate of $g_i = \frac{\phi_i}{\sum_{j=1}^{n}\phi_j}r$. More generally, if the set of sessions with queued data at time t is $S(t) \subseteq \{1,\ldots,n\}$, the session $i \in S(t)$ receives service at rate $\frac{\phi_i}{\sum_{j\in S(t)}\phi_j}r$ at time t.

In [10, 11], Parekh and Gallager presented a thorough examination of the GPS scheduling where the source traffic of each session is characterized as a (σ, ρ) *Linear Bounded Arrival Process* (LBAP) [2]. Given that the traffic of each session conforms to a LBAP (as would be the case when a session is regulated by a leaky bucket mechanism) and that the total arrival rate of all the sessions is smaller than the service rate, it was shown that the backlog and delay of each session are bounded from above for a broad class of GPS networks.

In [15, 17], GPS is studied when the traffic generated by sources is modeled by an *Exponentially Bounded Burstiness* (E.B.B.) process [14]. Under the E.B.B. model, performance bounds analogous to those of the deterministic model are obtained in [17]. Given that the appropriate stability conditions are satisfied, upper bounds on the backlog and delay tail distributions for each session sharing a single GPS server are obtained in the form of exponentially decaying functions and the departure process of each is shown to be an E.B.B process as well (the latter was also proved in [15]). In addition, a broad class of GPS networks with arbitrary topology is shown to be stable (see also [15]).

As we will see later, the difficulty in analyzing the GPS system in a stochastic setting arises because the instantaneous service rate of each session depends both on the arrival rate and queue length of all sessions. To circumvent this difficulty, [17] uses a decomposition approach to derive *upper bounds* on the session backlog and delay tail distributions. By generalizing of the notion of a *feasible ordering*, introduced in [10], an n-queue GPS system is decomposed in a way that the sessions are decoupled, so that they can be analyzed separately as n independent G/D/1/∞ queues.

2.2 Statistical Multiplexing System with Constant Service Process

In this section we briefly review the analysis of a single-queue statistical multiplexing system with Markov modulated fluid sources (MMFS) and a constant service process. We assume the capacity of the buffer is infinite and the rate of the service process is a constant c. Thus we have a MMFS/D/1/∞ queueing system. We model the input as a Markov modulated fluid source with state space $\mathcal{S}$, characterized by the pair $(\boldsymbol{M}, \boldsymbol{\lambda})$ with $\boldsymbol{M}$ being the generator of the Markov chain and $\boldsymbol{\lambda}$ the rate vector. Let $\overline{\lambda}$ denote the average rate. In order to have a stable system, we assume that $\overline{\lambda} < c$.

Let the random variables Σ and Q represent the stationary input state and the queue length, and define the stationary distribution vector $\pi(q) = \{\pi_s(q)\}_{s\in\mathcal{S}}$, where $\pi_s(q) = Pr\{\Sigma = s, Q \leq q\}$, $s \in \mathcal{S}, 0 \leq q < \infty$. Then the system can be described by the following Kolmogorov differential equation

$$\frac{d}{dq}\pi(q)\boldsymbol{D} = \pi(q)\boldsymbol{M}, \quad 0 \leq q < \infty \tag{1}$$

where $\boldsymbol{D} = \boldsymbol{\Lambda} - c\boldsymbol{I}$ ($\boldsymbol{I}$ is the identity matrix and $\boldsymbol{\Lambda} = \text{diag}(\boldsymbol{\lambda})$ is the system rate matrix) is the drift matrix.

The stationary queue length tail distribution $G(q) = Pr\{Q > q\} = 1 - \langle\pi(q), \boldsymbol{1}\rangle$ has the following spectral representation

$$G(q) = \sum_{i:\text{Re}(z_i)<0} a_i e^{z_i q}, \quad 0 \leq q < \infty \tag{2}$$

where $\{z_i\}$ are the solution of the generalized eigenvalue problem $z\phi\boldsymbol{D} = \phi\boldsymbol{M}$.

It is known [1] that there exists one real eigenvalue, z_1, called the *dominant eigenvalue*, that satisfies

$$0 > z_1 > \max_{i:\text{Re}(z_i)<0, i\neq 1}\{\text{Re}(z_i)\}. \tag{3}$$

From (2), it follows that

$$G(q) \approx e^{z_1 q} \quad q \to \infty. \tag{4}$$

Hence the dominant eigenvalue determines the asymptotic behavior of the queue length distribution.

We now turn our attention to the output process characterization. Although the aggregation of Markov modulated fluid sources is still a Markov modulated fluid source, the output process from the single-queue multiplexing system

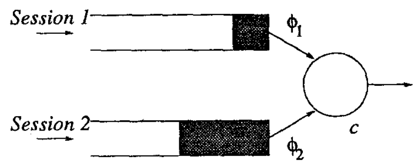

Figure 1: A two-queue GPS system.

is in general no longer so. A method to approximate the output process by a Markov modulated fluid source is proposed in [7], based on the techniques developed in [5]. The crucial idea is to aggregate the states $(\Sigma = s, Q = q)_{s \in \mathcal{S}, q>0}$ into a single busy state and to approximate the busy period by an exponential random variable.

Let $\mathcal{S}_U$ denote the set of underload states of the aggregate source, *i.e.*, $\mathcal{S}_U = \{s \in \mathcal{S} : \lambda_s < c\}$, and let s_b denote the busy state. The output process is approximately characterized by a Markov modulated fluid source $(\tilde{M}, \tilde{\lambda})$ on the state space $\tilde{\mathcal{S}} = \mathcal{S}_U \cup \{s_b\}$, where the generator matrix and the rate vector are given as follows:

$$\tilde{M} = \left[\begin{array}{c|c} M_{U,U} & a \\ \hline b & -(1/\bar{b}) \end{array}\right] \text{ and } \tilde{\lambda} = [\lambda_U, c]. \quad (5)$$

The matrix $M_{U,U}$ is a submatrix of M specifying the transition probability among the underload states. The vector a denotes the transition rates from the underload states to the busy period state s_b and the row vector b vice versa. Last $\bar{b}$ is the mean busy period. The bulk of output process characterization lies in computing a, b and $\bar{b}$, the details of which can be found in [7].

An important property of the above approximate characterization of the output process is that *all* of the stationary moments of the actual output process are identical to those of the process $(\tilde{M}, \tilde{\lambda})$.

3 Model

In this paper we focus on the two-queue GPS System (Figure 1). For $i = 1, 2$, the session i source is modeled as a Markov-modulated fluid process with an irreducible generator $M^{(i)}$ on state space $\mathcal{S}^{(i)} = \{1, \cdots, N^{(i)}\}$, and a rate vector $\lambda^{(i)} = \{\lambda_1^{(i)}, \cdots, \lambda_{N^{(i)}}^{(i)}\}$.

Let $\phi_i, i = 1, 2$ be the GPS assignment for the two sessions. Without loss of generality, we assume that $\phi_1 + \phi_2 = 1$. Hence each session is guaranteed a minimum service rate $g_i = \phi_i c$. For $i = 1, 2$, let $\bar{\lambda}_i$ denote the average input rate of session i source. As a necessary stability condition, we require that $\bar{\lambda}_1 + \bar{\lambda}_2 < c$.

For $i = 1, 2$, let $r_i(t) \in \{\lambda_1^{(i)}, \cdots, \lambda_{N^{(i)}}^{(i)}\}$ denote the rate of session i arrival process at time t, and for any $\tau < t$, define $A_i(\tau, t) = \int_\tau^t r_i(u)\,du$, the (cumulative) arrival process for session i. Similarly, let $s_i(t)$ denote the service/departure rate of session i at time t, and define $S_i(\tau, t) = \int_\tau^t s_i(u)\,du$, the session i departure process. Moreover, let $Q_i(t)$ denote the session i backlog at time t. Then, we have $Q_i(t) = \sup_{\tau \le t}\{A_i(\tau, t) - S_i(\tau, t)\}$.

The dynamics of the GPS system are described by the following sample path equations:

$$\frac{d}{dt}Q_1(t) = r_1(t) - \left[\phi_1 c + (\phi_2 c - r_2(t))^+ \mathbf{1}_{\{Q_2(t)=0\}}\right] \quad (6)$$

$$\frac{d}{dt}Q_2(t) = r_2(t) - \left[\phi_2 c + (\phi_1 c - r_1(t))^+ \mathbf{1}_{\{Q_1(t)=0\}}\right] \quad (7)$$

where $(x)^+ = \max\{x, 0\}$ and $\mathbf{1}_{\{.\}}$ is the indicator function. We note that from the definition of GPS, if both queues are not empty at time t, then $s_1(t) = g_1 = \phi_1 c$, $s_2(t) = g_2 = \phi_2 c$. If one of the sessions has an empty queue at time t, then the other session will receive the residual service from this session in addition to its guaranteed service rate.

4 Analysis

The difficulty in analyzing a GPS system lies in the fact that the service rate each session receives depends on the arrival rates and queue lengths of both sessions. As a consequence, an exact analysis of the queue length distributions of the two-queue GPS system involves solving a set of differential partial equations with coupled boundary conditions. This is very difficult to do.

In this section, we study the two-queue GPS system with Markov-Modulated Fluid Sources by combining the bounding approach in [17] and the spectral analysis approximation approach used in [6, 7]. As in [17], we decompose the two-queue GPS system into two independent queueing systems. By taking advantage of the specific structure of the arrival processes and the system, we are able to derive sample-path upper and lower bounds on the backlog process for each session. Using the approach in [6, 7], these sample-path relations can be converted to upper and lower bounds on the queue length distributions of the GPS system.

The rest of the section is organized as follows. In § 4.1, we study a single-queue statistical multiplexing system with a modulated service process using the techniques of [6, 13, 7], which lays the foundation for the lower and upper bound procedures proposed in § 4.2 and § 4.3. In § 4.2, we derive the lower bounds, and in § 4.3 we describe the procedure for obtaining the upper bounds. In § 4.4, the refined effective bandwidth approximation is discussed. In § 4.5, numerical and simulation results are compared.

4.1 Statistical Multiplexing System with Modulated Service Process

In this section we study a single-queue statistical multiplexing system with a modulated service process. The capacity of the buffer is assumed to be infinite. Arrivals are from a Markovian modulated fluid source characterized by a pair $(M^{(a)}, \lambda^{(a)})$ on state space $\mathcal{S}^{(a)}$. The service process is modulated by another Markovian modulated fluid process characterized by a pair $(M^{(s)}, \lambda^{(s)})$ on state space $\mathcal{S}^{(s)}$, which is independent of the arrival process. When the arrival process is in state $s^{(a)} \in \mathcal{S}^{(a)}$, it generates fluid at constant rate $\lambda_{s^{(a)}}^{(a)}$, while when the modulating service process is in state $s^{(s)} \in \mathcal{S}^{(s)}$, the queue is served at the rate $c - \lambda_{s^{(s)}}^{(s)}$, where c is the maximum service rate (hence $0 \le \lambda_{s^{(s)}}^{(s)} \le c$, $s^{(s)} \in \mathcal{S}^{(s)}$). We note this multiplexing system with modulated service process is very similar to the producer-consumer system studied in [9]. Let $Q(t)$ denote

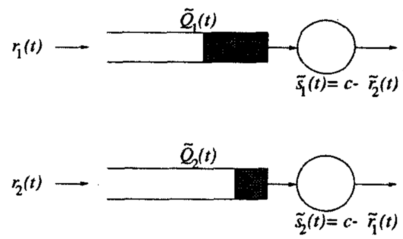

Figure 2: The decomposed lower bound system.

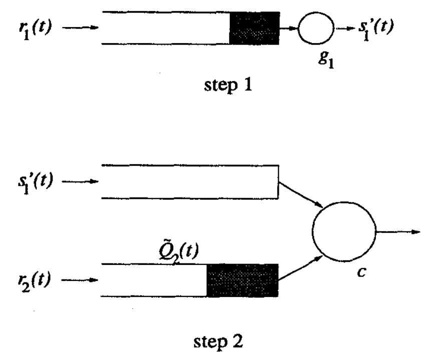

Figure 3: The two-step upper bound approach (*unbiased regime*).

the queue length of the system at time t; then $Q(t)$ satisfies the following sample path equation:

$$\frac{d}{dt}Q(t) = \lambda^{(a)}_{\Sigma^{(a)}(t)} - [c - \lambda^{(s)}_{\Sigma^{(s)}(t)}], \text{ if } Q(t) > 0, \quad (8)$$

where $(\Sigma^{(a)}(t), \Sigma^{(s)}(t))$ is a pair of random variables representing the states of the arrival process and the modulating service process at time t. By considering (8) it is clear that the system is equivalent to a statistical multiplexing system with constant service rate c, the input of which is provided by the superposition of the Markov modulated fluid sources $(\boldsymbol{M}^{(a)}, \boldsymbol{\lambda}^{(a)})$ and $(\boldsymbol{M}^{(s)}, \boldsymbol{\lambda}^{(s)})$ [13]. Thus, the analysis of the queue length distribution resembles the one we have briefly described in § 2.2; the stationary tail queue length distribution, $G(q) = Pr\{Q > q\}$ has the following spectral representation

$$G(q) = \sum_{i:\mathrm{Re}(z_i)<0} a_i e^{z_i q} \quad (0 \le q < \infty) \qquad (9)$$

where $\{z_i\}$ are the system eigenvalues.

4.2 Lower Bound for the Two-Queue GPS System

In this section, we derive the lower bound on the queue length distribution for each session.

We decompose the two-queue GPS system into two independent one-queue systems as depicted in Figure 2 where $r_i(t)$ is the arrival rate to the ith one-queue system at time t, $\tilde{s}_i(t)$ the service rate at time t, and $\tilde{Q}_i(t)$ the backlog of the queue at time t, $i = 1, 2$. We want to choose $\tilde{s}_i(t)$ such that $\tilde{Q}_i(t) \le Q_i(t)$ for all t.

Considering (6), if we assume that queue 2 is always empty, which amounts to ignoring the impact of the queue length of session 2 on the service rate of session 1, then queue 1 is effectively decoupled from queue 2. Hence, by choosing $\tilde{s}_1(t) = \phi_1 c + (\phi_2 c - r_2(t))^+ = c - \tilde{r}_2(t)$, where $\tilde{r}_2(t) = \min\{r_2(t), \phi_2 c\}$, we have an independent one-queue system for session 1 with a modulated service process described by $\tilde{s}_1(t) = c - \tilde{r}_2(t)$.

Similarly, by choosing $\tilde{s}_2(t) = \phi_2 c + (\phi_1 c - r_1(t))^+ = c - \tilde{r}_1(t)$, where $\tilde{r}_1(t) = \min\{r_1(t), \phi_1 c\}$, we have a one-queue system for session 2 with a modulated service process described by $\tilde{s}_2(t) = c - \tilde{r}_1(t)$.

The dynamics of the queue evolution of the two independent one-queue systems can be described as follows:

$$\frac{d\tilde{Q}_1(t)}{dt} = r_1(t) - [c - \tilde{r}_2(t)], \quad \text{if } \tilde{Q}_1(t) > 0 \quad (10)$$

$$\frac{d\tilde{Q}_2(t)}{dt} = r_2(t) - [c - \tilde{r}_1(t)] \quad \text{if } \tilde{Q}_2(t) > 0. \quad (11)$$

By comparing (6) and (7) with the above two equations, it is not too hard to see that for $i = 1, 2$, and for any t, the following holds

$$Q_i(t) \ge \tilde{Q}_i(t) \qquad (12)$$

Using the spectral analysis technique outlined in § 4.1, the stationary distribution of $\tilde{Q}_i(t)$ can be derived. For $i = 1, 2$, the input process for session i is characterized by the pair $(\boldsymbol{M}^{(i)}, \boldsymbol{\lambda}^{(i)})$ and the modulating processes by the pair $(\tilde{\boldsymbol{M}}^{(i)}, \tilde{\boldsymbol{\lambda}}^{(i)})$. According to the expressions for service rate, it follows that $\tilde{\boldsymbol{M}}^{(i)} = \boldsymbol{M}^{(i)}$ and $\tilde{\boldsymbol{\lambda}}^{(i)} = \left\{\tilde{\lambda}^{(i)}_1, \cdots, \tilde{\lambda}^{(i)}_{N^{(i)}}\right\}$ where $\tilde{\lambda}^{(i)}_s = \min\{\lambda^{(i)}_s, g_i\}, 1 \le s \le N^{(i)}$.

Given the stationary distribution of $\tilde{Q}_i$, from (12), we have that

$$\Pr\{\tilde{Q}_i > q\} \le \Pr\{Q_i > q\}, \ i = 1, 2, \ 0 \le q < \infty \quad (13)$$

where Q_i is the stationary version of $Q_i(t)$.

4.3 Upper Bound Approximation for the Two-Queue GPS System

In this section, we describe the procedure for obtaining the upper bound approximation to the queue length distribution for each session.

For ease of exposition, we consider two possible regimes separately:

1. $\overline{\lambda}_1 < g_1$ and $\overline{\lambda}_2 < g_2$;

2. $\overline{\lambda}_1 < g_1$ and $\overline{\lambda}_2 \ge g_2$, or $\overline{\lambda}_1 \ge g_1$ and $\overline{\lambda}_2 < g_2$.

We call the first regime *unbiased* and the second *biased*. In the biased regime, we assume that $\overline{\lambda}_1 < g_1$ and $\overline{\lambda}_2 \ge g_2$, without loss of generality. In this regime, the queue length of session 2 will typically not decrease unless the session 1 queue is empty. In this case, session 2 receives the residual service rate from session 1. The strict priority case considered in [7] is an example of the biased regime where $\phi_1 = 1$ and $\phi_2 = 0$.

We describe the procedure for obtaining an upper bound approximation for the unbiased regime first and, without loss of generality, we focus on session 2. To obtain the session 1 queue length distribution, we simply reverse the role of session 1 and session 2 below.

The procedure consists of two steps. In the first step, we characterize the departure process of session 1 while assuming session 2 is always busy. From (6), this is equivalent to consider a one-queue multiplexing system with a constant service rate g_1. Let $s'_1(t)$ denote the departure rate of this multiplexing system at time t. To characterize the departure process, we use the technique of [7] as described in §2.2 and approximate the departure process by a Markov Modulated Fluid Source with generator $M'^{(1)}$ and rate vector $\lambda'^{(1)}$. Given the departure process characterization of session 1, in the second step, we feed it to the two-queue GPS system to replace the corresponding arrival process, and then study the queue length distribution of session 2 (observe that since $s'_1(t) \leq g_1$ session 1 queue is always empty). The procedure is shown schematically in Figure 3. We observe that this new GPS system is equivalent to a one-queue multiplexing system with modulated service process characterized by $\tilde{s}_2(t) = c - s'_1(t)$. Let $\tilde{Q}_2(t)$ denote the queue length of this multiplexing system at time t. Then it can be shown [8] that for any t,

$$Q_2(t) \leq \tilde{Q}_2(t). \tag{14}$$

Therefore, the distribution of $\tilde{Q}_2(t)$ provides an upper bound on that of $Q_2(t)$, the queue length distribution of session 2 of the two-queue GPS system. As before, we focus on the stationary distributions only. Let $\tilde{Q}_2$ and Q_2 be the stationary versions of $\tilde{Q}_2(t)$ and $Q_2(t)$. Then from (14), we have

$$\Pr\{Q_2 > q\} \leq \Pr\{\tilde{Q}_2 > q\},\ 0 \leq q < \infty. \tag{15}$$

The distribution of $\tilde{Q}_2$ is computed using the spectral analysis technique outlined in § 4.1.

We now consider the biased regime. Recall that we have assumed that $\overline{\lambda}_1 < g_1 = \phi_1 c$ and $\overline{\lambda}_2 \geq g_2 = \phi_2 c$. Under this assumption, the upper bound approximation to the session 2 queue length process $Q_2(t)$ is obtained by using exactly the same two-step procedure for the unbiased regime outlined above. However, the upper bound approximation to the session 1 queue length process $Q_1(t)$ can be obtained in just one step by considering a one-queue multiplexing system with a service process of constant rate $g_1 = \phi_1$. This is due to the following: since $\overline{\lambda}_2 \geq g_2 = \phi_2 c$, the session 2 queue is very unlikely to be empty whenever session 1 is busy. This, in turn, implies that session 1 is most likely to receive only its minimum guaranteed service g_1 to clear its backlog. Hence its queue length process should be well approximated by that of the one-queue multiplexing system with a service process of constant rate g_1.

Observe that the procedure for obtaining the upper bound approximation for the biased regime also results in the decomposition of the two-queue GPS system into two independent one-queue multiplexing systems.

Last we remark that in case that $\phi_1 = 1, \phi_2 = 0$, the procedure for the biased regime reduces to that presented in [7] for the strict priority case.

We observe that the asymptotic decay rate of the upper bound approximation to the queue length distribution, obtained by means of the above procedures, equals the actual asymptotic decay rate as recently revealed in [16].

4.4 Effective Bandwidth Approximation

The lower and upper bound procedures proposed in the previous two sections require the complete analysis of the decomposed bounding system. As the dimension of the system grows, this becomes increasingly expensive computationally, especially when performed on-line. To remedy the situation, as in [7], we use the refined effective bandwidth approximation (REB). Namely, we approximate the queue length of a session by the following formula:

$$\Pr\{Q > q\} \approx Le^{-z^* q} \tag{16}$$

where z^* is the dominant eigenvalue of the system matrix in question and L is an appropriate prefactor.

As when $q = 0$, the approximation simply yields $G(0) \approx L$, thus L approximates the probability that the buffer is not empty. In [7], L is computed using the refined Chernoff large deviation approximation, which provides an upper bound on the stationary probability that the input rate exceeds the channel rate in a bufferless system. Here, we calculate L by simply computing the stationary probability that the input rate exceeds the channel capacity. This approximation underestimates the actual value of $G(0)$, as $G(0) \geq L$.

4.5 Numerical Investigation

In this section we investigate several examples to compare the proposed analytical bounds with simulation results.

For $i = 1, 2$, we assume session i is modeled by a superposition of K_i on-off fluid sources each characterized by the triple $(\alpha_i, \beta_i, \lambda_i)$. When the source is in the on state it generates fluid at a constant rate λ_i; when it is in the off state, it generates no traffic. The rate at which the source changes from the off-state to the on-state is α_i, the transition rate from the on-state to the off state is β_i. Session i has peak rate $K_i\lambda_i$ and average rate $K_i\frac{\lambda_i\alpha_i}{\alpha_i+\beta_i}$. The channel capacity is always assumed to be $c = 10.1$. We investigate how well the analytical bounds perform under different GPS assignments. The source parameters for both sessions are listed in Table 1. The average input rate for both session is $\overline{\lambda}_1 = \overline{\lambda}_2 = 2.857$, and the utilization factor of the channel is $\rho \approx 0.56$.

The analytical and simulation results for the following GPS assignments, $(\phi_1, \phi_2) = (0.7, 0.3)$ and $(0.8, 0.2)$ are shown in Figures 4 and 5, respectively. We note that the first assignment result in an *unbiased* GPS system, while the second one a *biased* system.

From the figures, we see that both lower and upper bounds are quite tight in all cases and the simulation values always lie between the upper bounds and the lower bounds. In particular, for small values of q, the bounds almost coincide with the simulation. As ϕ_1 increases from 0.7 to 0.8,

	K_i	α_i	β_i	λ_i
Session 1	10	0.4	1	1
Session 2	10	0.4	1	1

Table 1: System parameters for Example 1.

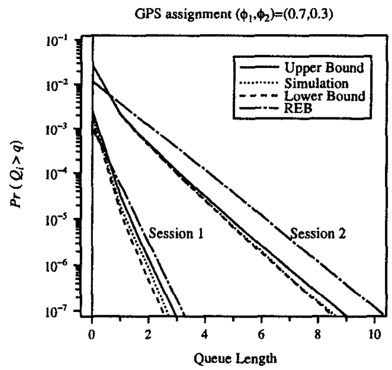

Figure 4: Bounds on queue length distributions for $(\phi_1, \phi_2) = (0.7, 0.3)$.

we obtain tighter upper and lower bounds. The reason for this behavior is explained in the next paragraph when we compare the mean busy and empty periods. The bounds computed using the refined effective bandwidth (REB) approximation underestimate the loss probability in the small buffer region, but are generally looser when the buffer size is large.

To further compare the analytical bounds with the simulation results, we look at the mean busy and empty periods from the analytical lower and upper bound systems (here empty period refers to an interval of time during which the queue is empty) and compare them with the simulated GPS system. The results are listed in Table 2. We see that both the analytical computed mean busy and empty periods are very close to the simulation values. In particular, we see that the lower bound system slightly underestimates the mean busy period while the upper bound system slightly overestimates it. The opposite occurs for the mean empty period. We observe that in all the cases the queue for each session is empty for an overwhelming fraction of time (the probability that $G(0) = \Pr\{Q_i > 0\}, i = 1, 2,$ *i.e.*, session i is busy, lies only in the range $0.0002 \sim 0.027$). It is not surprising that the probability of the system being busy is small considering the bursty sources and extreme low loss probability with the given buffer size.

From the description of the lower bound and upper bound systems in the previous sections, one can see that, when a session is in an empty period, its service rate is captured exactly, while when it is in a busy period, its service rate is either overestimated (in the case of the lower bound system) or underestimated (in the case of the upper bound system). As ϕ_1 increases and ϕ_2 decreases, session 1 is more likely to be idle while session 2 is more likely to be busy. Hence the approximation of the session 2 service rate during a busy period becomes more accurate. Moreover, as the service rate for session 1 becomes large, session 1 is less dependent on the behavior of session 2. Therefore, the analytical bounds tend to be more accurate. We also investigated how the busy period affects the quality of the bound; the results show that tight bounds are obtained even when the probabilities that the queue is not empty are quite high. Details can be found in [8].

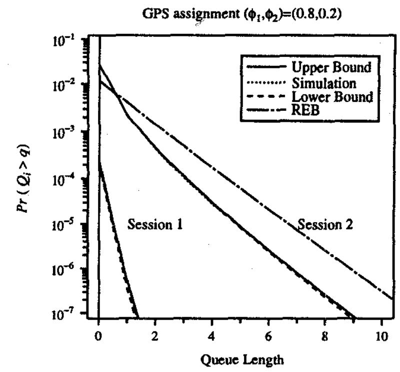

Figure 5: Bounds on queue length distributions for $(\phi_1, \phi_2) = (0.8, 0.2)$.

		$\phi_1 = 0.7$			$\phi_1 = 0.8$		
		Lower	Upper	Simul.	Lower	Upper	Simul.
Session 1	busy	0.34	0.40	0.36	0.30	0.32	0.31
	empty	206	156	203	1528	1346	1510
Session 2	busy	0.43	0.43	0.43	0.43	0.43	0.43
	empty	15.9	15.9	15.9	15.6	15.6	15.6

Table 2: Mean busy and empty periods.

5 Applications

In this section we investigate the issues in applying GPS scheduling to call admission control and bandwidth allocation and study the relative merits of GPS, strict priority and FIFO scheduling.

5.1 Loss Probability as the Sole QoS Requirement

Consider two classes of network traffic sharing a channel of capacity c, where each class has a desired loss probability bound η_i, $i = 1, 2$, as its QoS requirement. We are interested in comparing the three scheduling policies, GPS, strict priority and FIFO, in terms of the *admissible region* permitted by each policy, *i.e.*, the number of sources of each class admitted under the policy without violating the QoS requirement.

For simplicity, we model sources of both classes as homogeneous on-off sources characterized by a triple $(\alpha_i, \beta_i, \lambda_i)$, the meaning of which is exactly the same as in § 4.5. The parameters for the sources are listed in Table 3. The channel capacity c is 10.1.

Under GPS and strict priority scheduling, each class is assumed to have its own buffer of size q_i. We use

Class	α_i	β_i	λ_i	$\bar{\lambda}_i$	q_i	η_i
1	0.4	1	1	0.28	2	10^{-5}
2	0.5	0.5	1	0.5	1	10^{-7}

Table 3: System and on-off fluid source parameters for both classes.

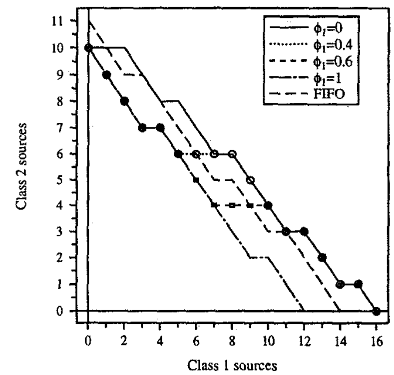

Figure 6: Admissible regions under GPS,priority and FIFO.

the stationary tail distribution $Pr\{Q_i > q_i\}$ in an infinite buffer system as a conservative estimate of the buffer loss probability for each class. Therefore, the loss QoS requirement can be represented as the constraint $Pr\{Q_i > q_i\} \leq \eta_i$. For FIFO scheduling, the two classes share a buffer of size $q = q_1 + q_2$. However, the loss QoS requirement will be determined by the more stringent one of the two classes, namely, we have the constraint that $Pr\{Q > q\} \leq \min\{\eta_1, \eta_2\}$, where Q is the stationary queue length process of aggregation of the two classes. As a first example, we assume that η_1 and η_2 are different. Values of η_i and q_i for each class are listed in Table 3. The admissible regions under the three scheduling policies are shown in Figure 6 where those under GPS are computed via the upper bound approximation procedure. From Figure 6, we see that under the GPS scheduling, the bigger ϕ_2 is, the larger the admissible region is, and thus the strict priority policy which gives the high priority to class 2 (corresponding to a GPS assignment $(\phi_1, \phi_2) = (0, 1)$) yields the largest admissible region, whereas the reverse strict priority (corresponding to a GPS assignment$(\phi_1, \phi_2) = (1, 0)$) the smallest. The admissible region under FIFO lies between the two strict priority cases.

In this example where the loss requirement of the two classes are considerably different, FIFO is not really appropriate, as the buffer loss requirement is determined by the more stringent one of the two, resulting in a smaller admissible region than the strict priority scheduling with class 2 as the high priority class. GPS scheduling with explicit bandwidth sharing (*i.e.*, $\phi_1 \neq 0$ and $\phi_2 \neq 0$) does not provide any benefit over the strict priority scheduling either. If we now consider the same two classes but with equal loss probability requirements, say, $\eta_1 = \eta_2 = 10^{-5}$,

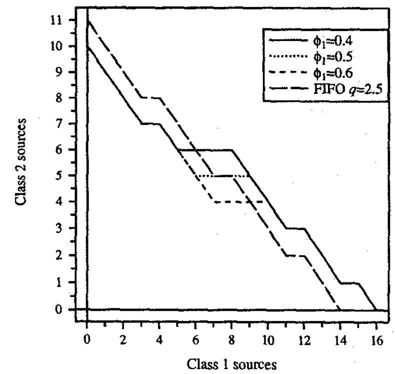

Figure 7: Admissible regions under GPS and FIFO.

FIFO gives the largest admissible region [8].

These examples indicate that, in the case that we have only two classes and that buffer loss probability is the sole QoS requirement, either strict priority or FIFO outperforms GPS scheduling with explicit bandwidth sharing. These observations are not surprising.

5.2 Loss Probability with Maximum Delay Constraint as QoS Requirement

In this section, we investigate how the three scheduling policies perform when we add an additional constraint on delay. Obviously delay is an important performance metric that should not be ignored when one or both of the classes are real-time. In the following example, we assume both classes have delay constraints; the observations made are also valid when only one of them has a delay constraint.

We consider two classes of real-time traffic sharing a channel of capacity c. For each class, in addition to the loss probability requirement, there is also a *maximum delay bound* $\hat{D}_i$ such that when traffic is not lost at the buffer, it must be serviced within $\hat{D}_i$ time units. In the case of GPS scheduling, suppose that the buffer size is q_i and $g_i = \phi_i c$ is the minimum guaranteed service rate, then $\frac{q_i}{g_i} \leq \hat{D}_i$ provides a sufficient guarantee that any real-time traffic that is not lost will not experience a delay more than $\hat{D}_i$. Thus, the maximum delay constraint $\hat{D}_i$ can be translated into a lower bound $\frac{q_i}{\hat{D}_i}$ on the minimum guaranteed service rate, or equivalently, $\phi_i \in [\frac{q_i}{\hat{D}_i c}, 1]$.

For illustration purposes, we assume that the source parameters for the two classes of real-time traffic and their loss requirements are the same as listed in Table 3, and add the delay constraints $\hat{D}_1 = 0.5$ for class 1 and $\hat{D}_2 = 0.25$ for class 2. As a result, we have $\phi_1 \in [0.4, 0.6]$ and $\phi_2 = 1 - \phi_1$. This effectively eliminates strict priority scheduling as a valid choice of scheduling policy. We now investigate the choice of ϕ_1 in this range that gives the largest admissible region under the given QoS constraints. The admissible regions computed via the upper bound method are shown in Figure 7 for three different GPS assignments $(\phi_1, \phi_2) \in \{(0.6, 0.4), (0.5, 0.5), (0.4, 0.6)\}$. Clearly, $(\phi_1, \phi_2) = (0.4, 0.6)$ gives the largest admissible region, which is expected as class 2 traffic has more *stringent* QoS requirement than class 1 traffic. In Figure 7 the admissible region under FIFO with buffer size $q = 2.5$ is also shown. Under FIFO, the loss probability requirement is $\eta_2 = 10^{-7}$, corresponding to the more stringent loss requirement of the two classes. The buffer size $q = 2.5$ is chosen so that the most stringent delay constraint $\hat{D}_2 = 0.25$ is always satisfied by the traffic that is not lost.

From the figure, we see that the admissible region is larger under FIFO when the number of class 1 sources is small, but smaller otherwise. Intuitively, under FIFO, when there are few class 1 sources, the shared buffer allows more class 2 sources to be admitted. On the other hand, under GPS, more class 1 sources can be admitted because of their less stringent QoS requirement.

This example indicates that GPS scheduling provides flexibility in providing explicit bandwidth sharing, which is absent in strict priority and FIFO scheduling, when QoS guarantees include both loss requirement and delay constraints.

6 Conclusion

In this paper, we studied GPS when arrivals are from Markov Modulated Fluid Sources. For simplicity, we focused on a two-class GPS system. We combined the bounding approach of [17] with the approximation approach based on spectral analysis techniques developed in [1, 5, 6, 7]. By employing the key idea of decomposition [17] and by taking advantage of the specific structure of MMFSs, we derived lower and upper bound systems that produce very accurate approximation to the original GPS system, as shown by numerical examples. Application of these bounds for GPS scheduling in the context of call admission control and bandwidth sharing was also illustrated, and comparison with FIFO and strict priority in different scenarios was made and discussed. We show that GPS scheduling provides flexibility in bandwidth allocation and sharing which is absent in strict priority and FIFO scheduling, while still satisfying the QoS requirement of each class. By providing a minimum guaranteed service rate to each class, delay constraints can be considered simultaneously with loss probability requirements. By removing the delay constraint used in our examples, our experience shows that strict priority is generally the best choice in providing loss probability guarantees when the loss probability requirements for the two classes are very different; whereas FIFO performs the best with two classes with very similar loss probability requirements. These observations should not be surprising.

Acknowledgement

The second author would like to thank Anwar Elwalid and Debasis Mitra for many inspiring conversations on the subject while he was a summer intern at AT&T Bell Lab, Murray Hill, New Jersey. The simulator we use in the paper is based on the simulator originally written by Anwar Elwalid. We would also like to thank Jim Kurose for many helpful comments.

References

[1] D. Anick, D. Mitra and M. M. Sondhi, Stochastic Theory of a Data Handling System with Multiple Sources, *Bell System Technical Journal*, Vol. 61, No. 8, Oct. 1982, pp. 1871-1894.

[2] R. L. Cruz, A Calculus for Network Delay, Part I: Network Elements in Isolation, *IEEE Transaction on Information Theory*, Vol. 37, No. 1, Jan. 1991, pp. 114-131.

[3] D. Clark, S. Shenker and L. Zhang, Supporting Real-Time Applications in an Integrated Service Packet Network: Architecture and Mechanism. In *Proceedings of ACM SIGCOMM'92*, pp. 14-26, 1992.

[4] A. Demers, S. Keshav and S. Shenker, Analysis and Simulation of a Fair Queueing Algorithm, *Journal of Internetworking: Research and Experience*, 1, pp. 3-26, 1990. Also in *Proceedings of ACM SIGCOMM '89*, pp. 3-12.

[5] A. Elwalid and D. Mitra, Analysis and Design of Rate-Based Congestion Control of High Speed Networks, I: Stochastic Fluid Models, Access Regulation, *Queueing Systems*, Vol. 9, 1993, pp. 29-64.

[6] A. Elwalid and D. Mitra, Effective Bandwidth of General Markovian Traffic Sources and Admission Control of High Speed Networks, *IEEE/ACM Trans. on Networking*, Vol. 1, No. 3, June 1993, pp. 329-357.

[7] A. Elwalid and D. Mitra, Analysis, Approximation and Admission Control of a Multi-Service Multiplexing System with Priorities, INFOCOM'95, 1995.

[8] F. Lo Presti, Z.-L. Zhang and D. Towsley, Bounds, Approximations and Applications for a Two-Queue GPS System, *Technical Report* 95-109, Computer Science Department, University of Massachusetts, December 1995.

[9] D. Mitra, Stochastic Theory of a Fluid Model of Producers and Consumers Coupled by a Buffer, *Adv. Applied Probability*, Vol. 20, 1988, pp. 646-676.

[10] 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 Transaction on Networking*, Vol. 1, No. 3, pp. 344-357, June 1993.

[11] A. K. Parekh and R. G. Gallager, A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks: The Multiple Node Case, *IEEE/ACM Transaction on Networking*, No. 2, Vol. 2, pp. 137-150, April 1994.

[12] S. Shenker, D. Clark and L. Zhang, A Scheduling Service Model and a Scheduling Architecture for an Integrated Services Packet Network, *Preprint*, 1993.

[13] T. Stern, A. Elwalid, Analysis of Separable Markov-Modulated Rate Models for Information-Handling Systems, *Adv. Applied Probability*, Vol. 23, pp. 105-139, 1991.

[14] O. Yaron, M. Sidi, Performance and Stability of Communication Networks via Robust Exponential Bounds, *IEEE/ACM Trans. on Networking*, Vol. 1, No. 3, pp. 372-385, 1993.

[15] O. Yaron, M. Sidi, Generalized Processor Sharing Networks with Exponentially Bounded Burstiness Arrivals, In *Proceedings of IEEE INFOCOM '94*, pp.628-634, June 1994.

[16] Z.-L. Zhang, Large Deviations and Generalized Processor Sharing: Upper and Lower Bounds, Part I Two-Queue Systems, *Technical Report* UM-CS-95-96, Computer Science Department, University of Massachusetts, November 1995.

[17] Z.-L. Zhang, D. Towsley and J. Kurose, Statistical Analysis of Generalized Processor Sharing Scheduling Discipline, *IEEE Journal on Selected Area in Communications*, Vol. 13, No. 6, pp. 1071-1080, August 1995.