



Call Admission Control Schemes under the Generalized Processor Sharing Scheduling

Zhi-Li Zhang, Zhen Liu, Jim Kurose, Don Towsley

► **To cite this version:**

Zhi-Li Zhang, Zhen Liu, Jim Kurose, Don Towsley. Call Admission Control Schemes under the Generalized Processor Sharing Scheduling. RR-2711, INRIA. 1995. <inria-00073980>

HAL Id: inria-00073980

<https://hal.inria.fr/inria-00073980>

Submitted on 24 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Call Admission Control Schemes under the Generalized Processor Sharing Scheduling

Zhi-Li Zhang, Zhen Liu, Jim Kurose, Don Towsley

N° 2711

Novembre 1995

PROGRAMME 1



*Rapport
de recherche*

Call Admission Control Schemes under the Generalized Processor Sharing Scheduling*

Zhi-Li Zhang, Zhen Liu, Jim Kurose, Don Towsley

Programme 1 — Architectures parallèles, bases de données, réseaux et systèmes distribués
Projet MISTRAL

Rapport de recherche n° 2711 — Novembre 1995 — 28 pages

Abstract: Provision of *Quality-of-Service* (QoS) guarantees is an important and challenging issue in the design of integrated-services packet networks. Call admission control is an integral part of the challenge and is closely related to other aspects of networks such as service models, scheduling disciplines, traffic characterization and QoS specification. In this paper we provide a *theoretical framework* within which call admission control schemes with multiple statistical QoS guarantees can be constructed for the Generalized Processor Sharing (GPS) scheduling discipline. Using this framework, we present several admission control schemes for both session-based and class-based service models. The theoretical framework is based on recent results in the statistical analysis of the GPS scheduling discipline and the theory of effective bandwidths. Both optimal schemes and suboptimal schemes requiring less computational effort are studied under these service models. The QoS metric considered is loss probability.

Key-words: Asymptotic decay rate, Call admission control, Effective bandwidths, Feasibility test, Generalized Processor Sharing scheduling, Loss probability, Network service models, Quality of service.

(Résumé : *tsvp*)

*Correspondence: Zhen LIU, INRIA Centre Sophia Antipolis, 2004 Route des Lucioles, B.P. 93, 06902 Sophia Antipolis, France. E-mail: liu@sophia.inria.fr

Mécanismes de Contrôle d'Admission sous la Discipline Processeur Partagé Généralisé

Résumé : La garantie de Qualité de Service (QS) est un aspect important dans la conception des réseaux à haut-débit. Le contrôle d'admission fait intégralement partie de ce défi et est étroitement lié aux autres aspects des réseaux tel que les modèles et disciplines de service, la caractérisation du trafic et la spécification de QS. Dans cet article nous fournissons un cadre théorique dans lequel des mécanismes de contrôle d'admission avec de multiple garanties statistiques de QS peuvent se construire pour la discipline de service processeur partagé généralisé (GPS). Ce cadre théorique se base sur les résultats récents en analyse probabiliste de la discipline GPS et sur la théorie de la bande passante effective. Nous présentons ainsi plusieurs mécanismes de contrôle d'admission pour les modèles de service tant au niveau client qu'au niveau session. Nous considérons le taux de perte comme critère de QS, et nous étudions des schémas optimaux et sous optimaux qui requièrent un coût de calcul moins important.

Mots-clé : taux de décroissance asymptotique, contrôle d'admission, bande passante effective, processeur partagé généralisé, taux de perte, qualité de service.

1 Introduction

Provision of *Quality-of-Service* (QoS) guarantees is an important and challenging issue in the design of integrated-services packet networks. Call admission control is an integral part of the problem. Clearly, without call admission control, providing QoS guarantees will be impossible. The task of call admission control can be most easily illustrated by considering the following question:

Given a new call/session that arrives to a network, can it be accepted by the network at its requested QoS, without violating existing QoS guarantees made to on-going calls?

This seemingly simple question turns out to be very complicated, as the issue of call admission control is closely related to other aspects of a network, such as service models, scheduling disciplines, traffic characterization and QoS specification. Call admission control with statistical QoS guarantees is a particularly important and challenging problem, particularly for a heterogeneous mixture of applications with differing QoS requirements.

In this paper we consider the call admission control issue for a network using Generalized Processor Sharing (GPS) scheduling at its switches and supporting multiple statistical QoS guarantees. We identify several service models that are likely to be provided by future integrated services packet networks, and propose corresponding call admission schemes for them. Included are both optimal schemes and suboptimal schemes requiring less computational effort. The theoretical foundation for our proposed schemes is based on recent results in the statistical analysis of GPS scheduling [27, 30] and the theory of effective bandwidths [17, 14, 15, 18, 12, 19, 25, 4, 16, 20]. The statistical QoS metric considered is loss probability.

We focus on GPS (also known as *Weighted Fair Queueing*) [10, 22, 23] because it provides controlled sharing of bandwidth and isolation among sessions (or classes). In [8, 24], GPS is recommended as a scheduling discipline where there are several different service classes which must be supported. The authors argue that perhaps the most important feature of GPS is its ability to isolate various service classes while, at the same time, allowing bandwidth sharing among classes. In addition, there is a rich literature analyzing GPS in a variety of settings. In [22, 23], per-session bounds on the worst-case backlog and delay are derived for both a single GPS server in isolation and a network of GPS servers under a deterministic setting. GPS has also been studied in the stochastic setting, where per-session bounds on backlog and delay tail distributions can be derived [27, 30]. These results make it possible to provide statistical QoS guarantees to applications with differing QoS requirements using GPS scheduling.

We use the theory of effective bandwidths because it provides the opportunity to place call admission control with multiple QoS requirements in a *formal and rigorous* framework. The theory of effective bandwidths has emerged recently as an elegant and promising approach to the problem of call admission control. It has been developed in the context of a single network switch or server with a finite capacity queue shared by many sessions, where the QoS metric in question is loss probability. Under this theory, a simple *asymptotically optimal* call admission control scheme exists for a network carrying a single class of traffic. However, this is not sufficient in an integrated services packet network where applications with quite different QoS requirements must co-exist. One of the contributions of this paper is to remedy this problem. For simplicity of exposition, the discussion will mostly focus on a single node case. Some of the schemes proposed can be extended to the end-to-end call admission control case in a fairly straightforward manner. The framework of our study is clearly theoretical. To the best of our knowledge, this is the first formal approach to the call admission control problem with multiple statistical QoS guarantees which accounts for various network service models. Even though many practical issues require resolution before these schemes can be applied in practice, we believe our study provides an important framework under which these issues can be investigated. For example, our schemes present a theoretical basis for addressing the call admission control issues for the service models (*e.g.*, the predictive service model) proposed in [8, 24].

The rest of the paper is organized as follows. Section 2 provides the necessary background for the understanding of the paper. This includes the notion of effective bandwidth in the context of the stochastic envelope process model [4] as well as the formal definition of GPS. Section 3 states a result on bounding the asymptotic decay rate of per-session backlog tail distribution and describes several feasibility tests based on this result. Section 4 presents several call admission control schemes with varying time-complexity for both session-based and class-based service models. Section 5 concludes the paper.

2 Preliminaries

2.1 Effective Bandwidths and Envelope Processes

The theory of effective bandwidths has been developed by many authors [17, 14, 15, 18, 12, 19, 25, 4, 16, 20] mostly in the context of a single-server queueing system with a queue shared by many sessions and where the QoS metric in question is queue loss probability. Intuitively, the effective bandwidth of a session is a quantity a^* associated with its arrival process that is equivalent to the service rate required to serve the session so that its QoS requirement can be satisfied asymptotically (*i.e.*, in the region of small loss




Figure 1: **Effective Bandwidth and QoS Requirement.**

probabilities). It lies between the peak rate and the average rate of the session and increases as the QoS requirement becomes more stringent (see Figure 1).

To introduce the concept of effective bandwidth in a more formal and general basis, we use C. S. Chang's notion of stochastic envelope processes (E.P.) [4] as a source traffic model. The E.P. model was originally defined for discrete time, but can be easily extended to continuous time. We will describe it below in the continuous time framework.

Consider the random rate process $\{a(t), t \geq 0\}$ which describes the actual traffic being offered to the network. Here $a(t)$ is the instantaneous arrival rate at time t . We assume that $a(t)$ is nonnegative and bounded for all $t \geq 0$. Then $A(\tau, t) = \int_{\tau}^t a(s) ds$ represents the cumulative arrivals over the time interval $[\tau, t]$ to the network. A will be referred to as an arrival process with rate process $\{a(t), t \geq 0\}$. For each $\theta \geq 0$, define

$$A^*(\theta, t) = \sup_{s \geq 0} \frac{1}{\theta} \log E e^{\theta A(s, s+t)}. \quad (1)$$

$A^*(\theta, t)$ is called the minimum envelope process of A with respect to θ in [4]. Associated with each $A^*(\theta, t)$ is the *minimum* envelope rate (MER) of A with respect to θ :

$$a^*(\theta) = \limsup_{t \rightarrow \infty} \frac{A^*(\theta, t)}{t}. \quad (2)$$

Given that the instantaneous arrival rate $a(t)$ is bounded for all $t \geq 0$, $a^*(\theta)$ is continuous and increasing in θ . Moreover, it can be shown [4] that $a^*(\theta)$ lies between the long term average rate and the long term peak rate.

The MER of an arrival process has the following *sub-additivity* property: let $\{a_i(t) : t \geq 0\}$, $1 \leq i \leq n$, be n independent rate processes, each with MER $a_i^*(\theta)$, and let $\{a(t) = \sum_{i=1}^n a_i(t) : t \geq 0\}$ be the aggregate rate process, then $a^*(\theta) \leq \sum_{i=1}^n a_i^*(\theta)$ where $a^*(\theta)$ is the MER of the aggregate rate process.

Now consider a G/D/1 queueing system with a server of constant service rate r and a queue of infinite capacity with n independent sessions, $\{a_i(t) : t \geq 0\}$, $1 \leq i \leq n$, sharing the queue. Let the service discipline be any work-conserving scheduling policy. Let $a_i^*(\theta)$ be the MER of session i , and $a^*(\theta)$ the MER of the aggregate rate process, $\{a(t) = \sum_{i=1}^n a_i(t), t \geq 0\}$. Then it can be proved [4] that the tail distribution of the queue length process $\{Q(t), t \geq 0\}$ satisfies the following relation

$$\lim_{q \rightarrow \infty} \frac{1}{q} \log \Pr\{Q(t) \geq q\} \leq -\theta \quad \text{if } a^*(\theta) < r \quad (3)$$

In particular, by sub-additivity of MER, (3) holds if $a^*(\theta) \leq \sum_{i=1}^n a_i^*(\theta) < r$.

The relationship in (3) can be made *tight* if stronger conditions on the rate processes $\{a_i(t), t \geq 0\}$ are imposed. For example, assume (see [4]) that $\{a_i(t) : t \geq 0\}$ is stationary and ergodic; $a_i^*(\theta) = \lim_{t \rightarrow \infty} \frac{A_i^*(\theta, t)}{t}$ exists for all $\theta \geq 0$; and $\theta a_i^*(\theta)$ is strictly convex and differentiable for $\theta \geq 0$. Under these assumptions, the MER is *additive*, i.e., $a^*(\theta) = \sum_{i=1}^n a_i^*(\theta)$, and the stationary queue length process Q satisfies

$$\lim_{q \rightarrow \infty} \frac{1}{q} \log \Pr\{Q \geq q\} \leq -\theta \quad \text{iff } a^*(\theta) = \sum_{i=1}^n a_i^*(\theta) < r. \quad (4)$$

In this case, the MER $a^*(\theta)$ is called the *effective bandwidth*. Its importance in network call admission control can be illustrated in the following example.

Suppose a network server has total bandwidth r and a shared queue of size q . The tail distribution $\Pr\{Q(t) \geq q\}$ of a G/D/1 ∞ system provides a conservative estimate of the loss probability of the network server. Let ϵ be a desired upper bound on the loss probability, then from (3) or (4), for q large enough, the condition $\sum_{i=1}^n a_i^*(\xi) < r$, where $\xi = \frac{-\log \epsilon}{q}$, implies that $\Pr\{Q \geq q\}$ can be approximately upper bounded by $e^{-\xi q} = \epsilon$. Clearly the test $\sum_{i=1}^n a_i^*(\xi) < r$ provides a basis for call admission control. Moreover, the *if and only if* relation in (4) indicates that this test is *asymptotically optimal*.

Let $\theta^* = \sup\{\theta \geq 0 : a^*(\theta)\}$. Equation (4) suggests that for q large, we have the following *effective bandwidth approximation* to the queue length tail distribution $Pr\{Q \geq q\}$:

$$Pr\{Q \geq q\} \approx e^{-\theta^* q}. \quad (5)$$

In other words, $a^*(\theta)$ *asymptotically* captures the bandwidth required so that $Pr\{Q \geq q\} \approx e^{-\theta q}$.

The effective bandwidth approximation to loss probability (5) is based on large queue asymptotics, and it may not be valid for small queue size. Moreover, it ignores the statistical multiplexing gain resulting from averaging over a large number of sources [6]. One improvement is to add an exponential prefactor Λ in front of $e^{-\theta^* q}$ in (5) (see, *e.g.*, [6, 13, 20]), where Λ can be viewed as an indication of the statistical multiplexing gains. In the same vein, more recent work has been conducted in investigating effective bandwidth approximations to loss probability when there are a large number of sources (see, *e.g.*, [9, 3, 11]). The class-based call admission control schemes proposed in this paper can be modified to incorporate this new approximation so that the statistical multiplexing gain due to large number of sources can be exploited.

2.2 Generalized Processor Sharing (GPS) Scheduling

Generalized Process Sharing 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 an infinitely divisible fluid (hence there is no notion of a “packet” in this traffic model [PG93a]). Assume that n sessions share 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 bandwidth share for 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 $S(t) \subseteq \{1, \dots, n\}$, session $i \in S(t)$ receives service at rate $\frac{\phi_i}{\sum_{j \in S(t)} \phi_j} r$ at time t .

The performance analysis of GPS scheduling have been carried out under both a deterministic setting and a stochastic setting. In [22, 23], Parekh and Gallager examined the GPS scheduling under Cruz’s (σ, ρ) -*Linear Bounded Arrival Process* (LBAP) model [7], where ρ captures the long-term rate of a session source and σ the size of the maximum burst of the source above the long term rate. Given that each session source 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 both in the case of a single GPS server in isolation and in the case of a broad class of GPS networks. Of particular interest are the so-called *Rate Proportional*

Processor Sharing (RPPS) GPS networks. For RPPS GPS networks, simple closed form expressions for bounds on the end-to-end delay of each session can be derived. In [27, 30], GPS scheduling was studied under the *Exponentially Bounded Burstiness* (E.B.B.) process model introduced in [26]. Here a source is a (ρ, Λ, α) -E.B.B. process if the probability that the amount of traffic generated by the source during $[\tau, t]$ exceeds $\rho(t - \tau) + \sigma$ is bounded above by $\Lambda e^{-\alpha\sigma}$ for all $\sigma \geq 0$, where ρ is called the long term *upper rate* of the arrival process, Λ the prefactor, and α the decay rate. Under the E.B.B. model, performance bounds analogous to those of the deterministic model can be obtained. 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. These bounds can be extended to a broad class of GPS networks. In particular, for RPPS GPS networks, the upper bounds on the backlog and delay tail distributions for each session have simple closed form expressions.

The aforementioned results, both deterministic and statistical, are derived via an important concept introduced by Parekh and Gallager: the notion of *feasible ordering*. Given the rates ρ_i (either the rate of an LBAP or an upper rate of an E.B.B. process), $1 \leq i \leq n$, an ordering of the sessions, s_1, s_2, \dots, s_n , is a feasible ordering with respect to $\{\rho_i\}_{1 \leq i \leq n}$ and $\{\phi_i\}_{1 \leq i \leq n}$ if for $i = 1, 2, \dots, n$ (note that by convention, $\sum_{j=1}^0 \rho_{s_j} = 0$),

$$\rho_{s_i} < \frac{\phi_{s_i}}{\sum_{j=i}^n \phi_{s_j}} \left(r - \sum_{j=1}^{i-1} \rho_{s_j} \right). \quad (6)$$

It can be shown that as long as $\sum_{i=1}^n \rho_i < r$, such a feasible ordering always exists [22].

A GPS assignment $\{\phi_i\}_{1 \leq i \leq n}$ is called an RPPS GPS assignment if for $1 \leq i \leq n$, $\rho_i < g_i = \frac{\phi_i}{\sum_{j=1}^n \phi_j} r$.

In particular, if $\rho_i = \phi_i$, then $\rho_i < g_i$ assuming that $\sum_{i=1}^n \rho_i < r$, hence the name *Rate Proportional Processor Sharing*. The notion of RPPS GPS assignment can be extended to a network of GPS servers in a straightforward manner. GPS scheduling with an RPPS GPS assignment is called RPPS GPS scheduling.

3 Bounds on Asymptotic Decay Rates and Feasibility Tests

The analytical results from [22, 23, 27, 30], which provide provable bounds on the backlog and delay tail distributions, provide a theoretical basis for performing call admission control using GPS scheduling. In the deterministic regime where worst-case deterministic QoS guarantees are provided, call admission control is relatively easy. On the other hand, in the stochastic regime where statistical QoS guarantees

are provided, call admission control appears to be considerably more complicated. One obvious question, for example, is what is the minimum service rate (or bandwidth) a session requires in order to satisfy its QoS guarantees? Here the notion of effective bandwidth proves to be a very natural indicator of service or bandwidth requirement. In the following sections, we show how the theory of effective bandwidths can be applied to address the call admission control issue under GPS scheduling.

3.1 Upper Bounds on Asymptotic Decay Rates

In order to apply the theory of effective bandwidths, we need to derive upper bounds on the asymptotic decay rate of the per-session backlog tail distribution, similar in form to (3) under GPS scheduling.

Suppose we have n sessions sharing a single GPS server with a given GPS assignment, $\{\phi_i\}_{1 \leq i \leq n}$. The session arrival processes are assumed to be independent. For session i , $1 \leq i \leq n$, $a_i^*(\theta)$ is the effective bandwidth function of its arrival process, well-defined for $\theta \geq 0$. Let r be the service rate of the GPS server, then a necessary stability condition is that $\sum_{i=1}^n a_i^*(0) < r$, *i.e.*, the sum of the (long term) average arrival rates of all the sessions cannot exceed the service rate. Given these assumptions, the following theorem can be shown to hold, the proof of which can be found in [29].

Theorem 1 *Under the assumptions stated in the preceding paragraph, we have that for each session i ,*

$$\limsup_{q \rightarrow \infty} \frac{1}{q} \Pr\{Q_i(t) \geq q\} \leq -\theta_i^* \quad (7)$$

where

$$\theta_i^* = \max_{F \subseteq N \setminus \{i\}} \sup\{\theta \geq 0 : a_i^*(\theta) < \frac{\phi_i}{\sum_{l \in F} \phi_l} (r - \sum_{l \in F} a_l^*(\theta))\}. \quad (8)$$

We remark that under some stronger conditions on the session arrival processes, tighter bounds on the asymptotic decay rate of the backlog tail distribution can be identified using the large deviation theory [28]. These bounds are, unfortunately, very difficult to compute.

3.2 Feasibility Tests

To answer the generic call admission control question posed earlier in Section , it is sufficient to answer the following *feasibility* question:

Given n sessions present in the system, each session having a given QoS guarantee, will the server be able to make the QoS guarantees for *all* the sessions?

In this section, we consider the asymptotic regime where the QoS requirements are expressed in terms of bounds on the asymptotic decay rates of the per-session backlog tail distributions. In other words, if $\epsilon_i(q)$ is the desired bound on the loss probability when the session i queue is of size q , *i.e.*, $\Pr\{Q_i(t) \geq q\} \leq \epsilon_i(q)$, then $\xi_i = \limsup_{q \rightarrow \infty} \frac{-\log \epsilon_i(q)}{q}$ is the desired bound on the asymptotic decay rate of session i backlog tail distribution. Thus ξ_i represents the QoS requirement of session i in our discussion.

From Theorem 1, we see that if we know θ_i^* for each session i , then the feasibility question raised above can be answered easily: if $\xi_i \leq \theta_i^*$ for all i , then the answer is *YES*; otherwise, *NO*. In the following we devise an efficient (*i.e.*, polynomial time in n) optimal feasibility test based on Theorem 1. By *optimality* here, we mean that if $\xi_i \leq \theta_i^*$, $1 \leq i \leq n$, then the test will output *YES*. Before we describe this test, we first introduce the notion of a *partial feasible partition*, which is an extension of the notion of a partial feasible ordering.

For any $F \subset N = \{1, 2, \dots, n\}$ with $m = |F| \leq n = |N|$, a partial feasible partition, F_1, F_2, \dots, F_m , of F with respect to θ is defined recursively as follows: for $1 \leq k \leq m$,

$$\gamma_k = \frac{1}{\sum_{j \in N \setminus F^{k-1}} \phi_j} (r - \sum_{j \in F^{k-1}} a_j^*(\theta)) \quad (9)$$

and

$$F_k = \{j \in F \setminus F^{k-1} : a_j^*(\theta) \leq \phi_j \gamma_k\}. \quad (10)$$

where $F^0 := \emptyset$ and $F^k := F_1 \cup \dots \cup F_k$ for $k \geq 1$.

We refer to the γ_k 's defined in (9) as the *associated delimiting numbers* for the partial feasible partition of F . We say F_k is well-defined if $F_k \neq \emptyset$. Suppose p is such that $F_p \neq \emptyset$ but $F_{p+1} = \emptyset$. Then $\beta_{p+1} = \dots = \beta_n$ and $F_{p+1} = \dots = F_m = \emptyset$ if $p < m$. It is possible that $F_1 = \emptyset$, or $F_m \neq \emptyset$. In the former case, the partial feasible partition with respect to θ is null. In the latter case, all F_k 's are singleton sets. Clearly, $F^p = F_1 \cup \dots \cup F_p \subseteq F$. The containment can be strict, hence the name *partial feasible partition* of F . In the case that $F = N$, $F^p = N$ if and only if $\sum_{j=1}^n a_j^*(\theta) \leq r$, in which case, F_1, \dots, F_p form a proper partition of N (called a *feasible partition* in [30]). Moreover, any ordering of the sessions in F_1 followed by any ordering of the sessions in F_2 , *etc.*, will produce a feasible ordering of the n sessions.

The notion of a partial feasible partition captures the inherent “priority” among the sessions and how the dynamic sharing of bandwidth among the sessions under GPS scheduling determines the decay rate of the backlog tail distribution of each session. For example, let F_1, \dots, F_n be the partial feasible partition of N with respect to a given θ . For a session i in F_k , $1 \leq k \leq n$, its backlog will decay at a rate of at least θ when the queues of sessions in F_l , $1 \leq l < k$, become empty, thus session i having a service rate of at least $\phi_i \gamma_k$. In particular, for a session in F_1 (if it is not empty), the minimum guaranteed service rate g_i suffices to ensure that the backlog of session i decays at a rate of at least θ .

A partial feasible partition and its associated delimiting numbers exhibit the following monotonicity properties, the proof of which is relegated to Appendix A.

Lemma 2 *For any $F \subseteq N$ with $m = |F|$, let F_1, F_2, \dots, F_m be the partial feasible partition of F with respect to θ and $\gamma_1, \gamma_2, \dots, \gamma_m$ be the associated delimiting numbers. We have*

(a) $\gamma_1 \leq \gamma_2 \leq \dots \leq \gamma_m$.

(b) *For any $E \subseteq F$, let E_1, E_2, \dots, E_l , $l = |E|$, be the partial feasible partition of E with respect to θ and $\eta_1, \eta_2, \dots, \eta_l$ be the associated delimiting numbers. Then $\eta_k \leq \gamma_k$ and $E^k := E_1 \cup \dots \cup E_k \subseteq F^k := F_1 \cup \dots \cup F_k$, $1 \leq k \leq l$.*

We now describe the optimal feasibility test: the test outputs *YES* if and only if $\xi_i \leq \theta_i^*$ for all i , $1 \leq i \leq n$, where ξ_i is the session i QoS requirement defined in the beginning of this section and θ_i^* is defined in (8).

For $1 \leq i \leq n$, let $H_{i,1}, \dots, H_{i,n}$ be the partial feasible partition of N with respect to ξ_i , and $\beta_{i,1}, \dots, \beta_{i,n}$ the corresponding set of delimiting numbers. The test is described in pseudo-code as follows:

Test 1 : Optimal Feasibility Test:

```

for  $i := 1$  to  $n$  do
  if  $i \notin H_i^n := \cup_{j=1}^n H_{i,j}$ 
    then output NO and stop;
endfor;
output YES and stop.

```

To see why this test gives the correct answer, observe that, if $i \in H_i^n$, then there exists l , $1 \leq l \leq n$, such that $H_{i,l} \neq \emptyset$ and $i \in H_{i,l}$. Hence by Theorem 1, $\xi_i \leq \theta_i^*$. We also claim that if $\xi_i \leq \theta_i^*$, then $i \in H_i^n$. Therefore the test is optimal. The claim is shown as follows. From $\xi_i \leq \theta_i^*$, we see that, from Theorem 1

and the increasingness property of a^* , there exists an $F \subseteq N_i$ such that

$$a_i^*(\xi_i) < \frac{\phi_i}{\sum_{j \notin F} \phi_j} (r - \sum_{j \in F} a_j^*(\xi_i)). \quad (11)$$

Let $F' = F \cup \{i\}$ and $m = |F'|$, thus $1 \leq m \leq n$. Consider the partial feasible partition, F_1, \dots, F_m , of F' and the partial feasible partition, H_1, \dots, H_n , of N with respect to ξ_i . Let $\gamma_1, \dots, \gamma_m$ be the associated delimiting numbers for the F' partition and β_1, \dots, β_n for the N partition. As $F' \subseteq N$, from Lemma 2(b), we have that $\gamma_k \leq \beta_k$ and $F^k \subseteq H^k$, $1 \leq k \leq m$. It can be shown that there exists k^* , $1 \leq k^* \leq m$ such that $i \in F_{k^*}$, as otherwise it contradicts (11), hence $i \in F^m$. But as $F^m \subseteq H^m \subseteq H^n$, we have $i \in H^n$. This establishes the optimality of the feasibility test.

We remark that for each i , $H_{i,1}, \dots, H_{i,n}$ can be constructed efficiently by using a sorted list of $\frac{a_j^*(\xi_i)}{\phi_j}$, $1 \leq j \leq n$. This takes $O(n \log n)$ time for each i , and $O(n^2 \log n)$ time in total. Once $H_{i,1}, \dots, H_{i,n}$ is given, the rest of the test can be carried out in $O(n^2)$ time. Hence, our optimal feasibility test takes $O(n^2 \log n)$ time. In many circumstances, a faster and simpler feasibility test is more desirable, although such a test may be “sub-optimal” in the sense that it rejects calls that could actually be admitted. Since the complexity of the optimal feasibility test mostly lies in the construction of the partial feasible partitions, simpler but suboptimal tests can be derived by using only part of this information. We provide two such examples.

Test 2 RPPS Feasibility Test:

```

for  $i := 1$  to  $n$  do
  if  $a_i^*(\xi_i) > g_i = \beta_{i,1}$ 
  then output NO and stop;
endfor;
output YES and stop.

```

This test is equivalent to checking whether $i \in H_{i,1}$ for all i and does not require that the partial feasibility partition be constructed. Clearly, this test takes only $O(n)$ time. Note that if we regard $\rho_i = a_i^*(\xi_i)$ as the “rates” of the session, since $\rho_i \leq g_i$, the sessions admitted are scheduled according to a RPPS-like GPS policy: each session is guaranteed a minimum bandwidth independent of other sessions. Hence we call this the RPPS feasibility test.

A slightly more sophisticated test which subsumes the RPPS feasibility test is described below. For each i , let $I_i = \{j \neq i : a_j^*(\xi_i) \leq \phi_j \beta_{i,1}\}$ and define $\nu_i = \frac{1}{\sum_{j \in N \setminus I_i} \phi_j} (r - \sum_{j \in I_i} a_j^*(\xi_i))$.

Test 3 Idle-Set Feasibility Test:

for $i := 1$ to n do
 if $a_i^*(\xi_i) > \phi_i \nu_i$
 then output *NO* and stop;
endfor;
output *YES* and stop.

Note that since $\nu_i \geq \frac{1}{\sum_{j \in N \setminus I_i} \phi_j} (r - \sum_{j \in I} \phi_j \beta_{i,1}) = \beta_{i,1}$, $e_i(\xi_i) \leq \phi_i \nu_i$ if and only if $i \in H_i^2 := H_{i,1} \cup H_{i,2}$.

Hence this test essentially uses the first two sets of the partial feasible partitions. The test is called the *Idle-Set Feasibility Test* because I_i contains the sessions that will likely be idle as $a_j^*(\xi_i) \leq \phi_j \beta_{i,1} = g_j$ for $j \in I_i$. Thus, session i will receive a service rate of at least $\phi_i \nu_i$ most of the time. The sets I_1, \dots, I_n can be constructed in $O(n \log n)$ by observing that $I_j \subseteq I_i$ if $\xi_j \leq \xi_i$, hence sorting ξ_1, \dots, ξ_n in increasing order, and for each j , a binary search of $\phi_j \beta_{i,1} = g_j$ in the sorted list $a_j^*(\xi_1), \dots, a_j^*(\xi_n)$ will decide which I_i to place j . This takes $O(n \log n)$ time, and once I_1, \dots, I_n is known, the rest of the test can be done in linear time.

As a comparison, we look at another simple test which is not unique to the GPS scheduling.

Test 4 The Aggregate Feasibility Test:

$\hat{\xi} := \max_{1 \leq i \leq n} \xi_i$;
if $\sum_{j=1}^n a_j^*(\hat{\xi}) \leq r$
then output *YES* and stop;
else output *NO* and stop.

The above condition $\sum_{j=1}^n a_j^*(\hat{\xi}) < r$ is equivalent to $\sum_{j=1}^n a_j^*(\xi_i) \leq r$ for all i . The condition ensures that ξ_i is a bound on the asymptotic decay rate of the aggregate backlog distribution of all the queues. It is thus also a bound on each individual queue. In other words, this test is oblivious of the queue scheduling policy as long as it is work-conserving, so it is applicable to such queue policies as GPS, priority-queue and head-of-line scheduling. Clearly, the test only takes $O(n)$ time.

It is not difficult to construct scenarios where Test 2 or Test 3 returns *YES* and Test 4 does not, or Test 4 returns *YES* but Test 2 and Test 3 return *NO*.

Class	α_i	β_i	λ_i	$\bar{\lambda}_i$	q_i	ϵ_i
1	0.025	0.045	1	0.357	100	10^{-3}
2	0.5	0.5	2	1	10	10^{-9}

Table 1: System and On-Off Fluid Source Parameters for Both Classes

3.3 A Numerical Example

In this section we present a simple numerical example that illustrates the use of feasibility tests described in the previous section. We consider a GPS server with two classes of traffic, each with its own queue. The rate of the server is r . For class i , $i = 1, 2$, the GPS assignment is ϕ_i , with $\phi_1 + \phi_2 = 1$, the queue size for class i is q_i , and the loss probability requirement is ϵ_i , *i.e.*, the probability of loss due to queue i overflow should be bounded from above by ϵ_i . There are n_i sources with each class i . Sources belonging to the same class are identical and modeled by the on-off fluid source characterized by a triple $(\alpha_i, \beta_i, \lambda_i)$. When a class i source is in the on-state, it generates traffic in 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 rate from the on-state into the off-state is β_i . Hence the peak rate of the source is λ_i and the average rate, $\bar{\lambda}_i$, is $\frac{\lambda_i \alpha_i}{\alpha_i + \beta_i}$. The effective bandwidth function for a single class i source is given by the following expression [1]:

$$a_i^*(\theta) = \frac{\theta \lambda_i - \alpha_i - \beta_i + \sqrt{(\theta \lambda_i - \beta_i + \alpha_i)^2 + 4 \alpha_i \beta_i}}{2\theta}. \quad (12)$$

By the additivity of effective bandwidth function, if there are n_i sources in class i , then the aggregate effective bandwidth function for the class is $n_i a_i^*(\theta)$.

We are interested in comparing the feasible regions under the feasibility tests described in the previous section. In other words, we look at the number of class 1 and class 2 sources that can be admitted into the system under those tests such that the designated QoS requirements for both queues are satisfied. Note that, since we have only two classes, the idle-set feasibility test is identical to the optimal feasibility test.

The source and system parameters for both classes are listed in Table 1. The service rate is $r = 100$. Note that the class 1 source is burstier and has a smaller average rate than class 2 source. Moreover, class 1 also has a less stringent QoS requirement, since $\xi_1 = \frac{-\log \epsilon_1}{q_1} \approx 0.069 < \xi_2 = \frac{-\log \epsilon_2}{q_2} \approx 2.07$.

We look at two scenarios: first one with a GPS assignment $\phi_1 = \phi_2 = 0.5$; second one with a GPS assignment $\phi_1 = 0.3$ and $\phi_2 = 0.7$. The results are shown in Figure 2.