Title	GPS schedulers and gaussian traffic
Author(s)	Mannersalo, Petteri; Norros, Ilkka
Citation	21st Annual Joint Conference of the IEEE Computer and Communications Societies, IEEE INFOCOM 2002. New York, USA, 23 - 27 June 2002, pp.1660-1667
Date	2002
URL	http://dx.doi.org/10.1109/INFCOM.2002.1019419
Rights	Copyright © (2002) IEEE. Reprinted from 21st Annual Joint Conference of the IEEE Computer and Communications Societies, IEEE INFOCOM 2002 . This article may be downloaded for personal use only

GPS Schedulers and Gaussian Traffic

Petteri Mannersalo, Ilkka Norros

***Abstract*— This article considers Gaussian flows which are fed into a GPS (*Generalized Processor Sharing*) scheduler. The system is analyzed using a most probable path approach. This method gives quite good approximations for performance measures, like queue length distributions in the full range of queue levels. The approximations are based on the distinction whether it is more probable that an aggregated queue consists of traffic from one class only or whether it is a combination of several classes. The approximate queue length distribution for a specific flow is then calculated either using the Empty Buffer Approximation or the authors' Rough Full Link Approximation, respectively.**

I. Introduction

TRAFFIC carried in a modern network is extremely complex and flows are usually controlled by feedback loops. Thus it is often fruitless — and also impossible — to use traditional performance analysis of simple queues. When observing aggregated traffic (aggregation both in time and in number of flows), Gaussian character seems to be always appearing when the level of aggregation is large enough [1]. The reason for that is, of course, the *Central Limit Theorem* which guarantees that summing independent variables with finite variance is going to converge to a Gaussian random variable.

The amount of aggregation which is needed in order to accept the Gaussian assumption varies a lot depending on the specific network environment. The appearance of Gaussian distributions is seen clearly in access networks, like modem pools, where access lines have limited capacities [2]. On the other hand, in Differentiated Services networks, traffic management operations are performed only on the level of traffic aggregates, not for individual flows (see e.g. [3]). Gaussian models might be justified in these both cases.

Service guarantees are playing an ever increasing role. This calls for traffic differentiation. A basic mechanism for differentiation is priority. Besides total precedence of one class before another, there are softer variants like drop precedences and guaranteed minimal throughput for each class. The *Generalized processor sharing* (GPS) discipline is a theoretical model which isolates traffic flows and provides service differentiation. Both deterministic and stochastic bounds (see e.g. [4], [5], [6]) have been derived for GPS systems. Moreover, large deviation based approaches have been used to found asymptotical decay rates of different performance measures (see e.g. [7], [8], [9], [10], [11]).

The requirement that incoming traffic is infinitely divisible makes the GPS scheduling policy unrealistic. However, there are many practical implementations for packetized traffic (e.g., packet by packet GPS and virtual clock scheduling). Usually, the difference between a real-life implementation and the theoretical GPS is not very large. For the packet-by-packet GPS schedulers, for example, it is easy to show that the extra delay due to indivisibility of packets is never larger than the server time of maximum sized packet served at the full server rate (see [4], [12]).

Petteri Mannersalo and Ilkka Norros are both with VTT Information Technology, P.O. Box 1202, FIN-02044 VTT, Finland. E-mail: {ilkka.norros,petteri.mannersalo}@vtt.fi

This paper is a continuation to a series of studies on queues with Gaussian input. The first papers studied a queue with fractional Brownian motion (FBM) input through a trivial lower bound in [13] and applying the generalized Schilder's theorem in [14]. The latter approach was extended to ordinary queues with general Gaussian input in [15], [16], and further to priority queues in [17]. Here we apply a similar machinery to queues served by a GPS scheduler.

Our main goal is to introduce approximations based on the most probable paths for the queue length distributions of a Gaussian GPS queueing system. Moreover, we try to emphasize that this approach is easily implemented as an expert tool which can be used in dimensioning without comprehensive knowledge of the mathematics of Gaussian processes. As a by-product of the examples, we show that the simple-minded mean rate based resource sharing does not work with Gaussian traffic.

The paper is structured as follows. GPS schedulers, Gaussian traffic and the most probable paths are defined in Section II. Approximations of the queue length distributions are introduced in Section III. In Section IV, we show how the most probable path approach can be applied in a specific GPS example. Finally, some conclusions are drawn in Section V.

II. GPS queueing system

A. GPS scheduler

Generalized Processor Sharing is a work-conserving scheduling discipline defined under the assumption that sources are described by fluid models. We consider k input classes sharing a common server of capacity c. For each flow $i = 1, \dots, k$, a guaranteed server rate $\mu_i c$ is assigned. We assume $\mu_i > 0$, $i = 1, \dots, k$, and

$$\sum_{i=i}^{k} \mu_i = 1.$$

Let $S^{\{i\}}(s,t)$ be the amount class i traffic served in a time interval $[s,t]$. Formally, the GPS scheduler guarantee can be written as follows (see [4]):

$$\frac{S^{\{i\}}(s,t)}{S^{\{j\}}(s,t)} \geq \frac{\mu_i}{\mu_j} \quad j = 1, \dots, k \tag{1}$$

for all classes i which are *backlogged throughout the interval* $[s,t]$.

B. Gaussian input traffic and queueing

Denote the cumulative arrival process of class i by $A^{\{i\}} = \{A_t^{\{i\}}\}_{t\in\mathbb{R}}$ and the amount of class i arriving traffic in an interval $(s,t]$ by $A^{\{i\}}(s,t) = A_t^{\{i\}} - A_s^{\{i\}}$. For the superposition of classes i and j we write $A_t^{\{i,j\}} = A_t^{\{i\}} + A_t^{\{j\}}$, and the total input is denoted as $A_t = \sum_{i=1}^k A_t^{\{i\}}$.

Assume that the arrival processes $A^{\{i\}}$ are *independent continuous* Gaussian processes with *stationary increments*. These processes can be parameterized as

$$A_t^{\{i\}} = m_i t + Z_t^{\{i\}},$$

where m_i is the mean input rate and $Z^{\{i\}}$ is a centered continuous Gaussian process with variance $v_i(t) = \operatorname{Var} Z_t^{\{i\}}$ and (necessarily continuous) covariance function

$$\Gamma_i(s,t) = \operatorname{Cov}(Z_s^{\{i\}}, Z_t^{\{i\}}) = \frac{1}{2}(v_i(s) + v_i(t) - v_i(s-t)).$$

Remark: From the point of view of queueing theory, the assumption of Gaussian input is never fully acceptable. There is always a positive probability of negative input, which is nonsense from practical point of view and destroys many classical arguments on the theoretical side.

An elegant definition which results in positive queue length processes even in the Gaussian case was given by Massoulie [11]. In a GPS system with unlimited buffers, the queue of class i, $Q_t^{\{i\}}$, and the total queue $Q_t = \sum_{i=1}^k Q_t^{\{i\}}$ satisfy

$$Q_t^{\{i\}} = \sup_{s\le t}(A^{\{i\}}(s,t) - \mu_i c T(s,t)) \tag{2}$$

$$Q_t = \sup_{s\le t}(\sum_{i=1}^k A^{\{i\}}(s,t) - c(t-s)), \tag{3}$$

where $T(s,t) = T_t - T_s$ and T_t is a non-decreasing stochastic process with $T_0 \equiv 0$. Thus, $\mu_i c T(s,t)$ presents in a sense the amount of *potential service* for each class i in time interval $(s,t]$. If all the classes are queueing on interval $[s,t]$, then $T(s,t) = t-s$, i.e., everyone gets exactly its guaranteed service. Otherwise, we require that $T(s,t) \ge (t-s)$ and its role is to redistribute the excess capacity. Note that a negative input is considered as an extra service capacity. The equations (2) and (3) uniquely define the $k+1$ queueing processes $Q^{\{1\}}, \ldots, Q^{\{k\}}, Q$ (see [11]). In order to define the time change T uniquely, one should add an extra condition, like T being the smallest process satisfying (2) and (3) with $T(s,t) \ge t-s$.

It is easy to show that these equations define a GPS queueing system, i.e., if processes $Q^{\{1\}}, \ldots, Q^{\{k\}}$ satisfy equations (2) and (3), then the corresponding input processes $A^{\{1\}}, \ldots, A^{\{k\}}$ are served according to the GPS policy: If the queue of class i traffic is non-empty throughout interval $[s,t]$, then the total amount of service it gets during that period is $S^{\{i\}}(s,t) = \mu_i c T(s,t)$. On the other hand, this same number is the maximum service a class could get; backlogged or not. Thus we get inequality (1). Equation (3) takes care of the work conservation.

If input traffic of some class has a continuous rate process, i.e., it is differentiable, and its rate is smaller than the guaranteed minimum rate, then it should not queue at all. This is one of the principal results for GPS queues and has been proved long ago. However, in order to demonstrate the power of the Massoulie's definition and to show that a negative input does not cause problems, we decided to write a proof.

Lemma 1: Consider the queueing system defined by (2) and (3). If a input process $A_t^{\{i\}}$ is differentiable, $\frac{d}{dt}A_t^{\{i\}} \le \mu_i c$ on some interval $[t_1,t_2]$, and $Q_{t_1}^{\{i\}} = 0$, then $Q_t^{\{i\}} = 0$ for all $t \in [t_1,t_2]$.

Proof: By the assumption

$$\frac{A^{\{i\}}(s,t)}{t-s} \le \mu_i c \quad \text{if } t_1 \le s \le t \le t_2.$$

Since $T(s,t) \ge t-s$ and $Q_{t_1}^{\{i\}} = 0$, we have for any $t \in [t_1,t_2]$

$$\begin{aligned} Q_t^{\{i\}} &= \sup_{t_1\le s\le t}\left(A^{\{i\}}(s,t) - \mu_i c T(s,t)\right) \\ &= \sup_{t_1\le s\le t}\left((t-s)\frac{A^{\{i\}}(s,t)}{t-s} - \mu_i c T(s,t)\right) \\ &\le \mu_i c \sup_{t_1\le s\le t}((t-s) - T(s,t)) = 0. \end{aligned}$$

■

Gaussian processes with stationary increments are not smooth if $\lim_{t\to 0} v(t)/t^2 = \infty$. For example, realizations of fractional Brownian motion are almost surely nowhere differentiable. On the other hand, the most probable paths (defined in the next subsection) are usually differentiable almost everywhere.

In teletraffic, the fractional Brownian motion model (see [18]) has formed a focus of study for a while. But there is no need to restrict to that small class of Gaussian processes. By similar arguments as in [17], it can be shown that any Gaussian process with stationary increments whose variance function satisfies the asymptotic condition (4) below is feasible.

Proposition 1: Storage processes $Q^{\{i\}}$, $i = 1,\ldots,k$, are finite almost surely if $\sum_{i=1}^k m_i < c$ and there exists $\alpha < 2$ such that

$$\lim_{t\to\infty}\frac{v_i(t)}{t^\alpha} = 0, \quad i = 1,\ldots,k. \tag{4}$$

Remark: An arbitrary function satisfying condition (4) is not necessarily a valid variance function for a Gaussian process; the corresponding covariance function must be positively semidefinite.

C. Most probable paths and large deviations

A large deviation principle for Gaussian measures in Banach space is given by the generalized Schilder's theorem (Bahadur and Zabell [19], see also [20], [21]). Its use to the study of queueing systems was introduced in the case of fractional Brownian motion in [14] and extended to arbitrary Gaussian traffic in [15], [16] and to priority queues in [17]. The present paper applies it to processor sharing schedulers.

Let us consider a Banach space

$$\Omega = \left\{\omega : \omega \text{ is continuous } \mathbb{R}\to\mathbb{R},\ \omega(0)=0,\ \lim_{t\to\infty}\frac{\omega(t)}{1+|t|} = \lim_{t\to-\infty}\frac{\omega(t)}{1+|t|} = 0\right\}$$

equipped with a norm

$$\|\omega\|_\Omega = \sup\left\{\frac{\omega(t)}{1+|t|} : t \in \mathbb{R}\right\}.$$

Let P be the unique probability measure on the Borel sets of $\Omega^k = \Omega \times \cdots \times \Omega$ such that random variables $Z_t^{\{i\}}(\omega_1, \ldots \omega_k) = \omega_i(t)$ form independent centered Gaussian processes with covariance functions $\Gamma_i(\cdot,\cdot)$. Condition (4) guarantees that $\lim_{t\to\pm\infty} Z_t^{\{i\}}/|t| \to 0$ a.s. – see [16].

Furthermore, let R be the reproducing kernel Hilbert space of $Z = (Z^{\{1\}}, \ldots, Z^{\{k\}})$ (for more information about reproducing kernel Hilbert spaces see e.g. [22]). In our case of independent processes $Z^{\{i\}}$, R is a Hilbert space spanned by the covariance functions $(\Gamma_1(t_1,\cdot), \ldots, \Gamma_n(t_k,\cdot))$, $t_i \in \mathbb{R}$, and having the properties

$$\begin{aligned}\langle(\Gamma_1(t_1,\cdot),\ldots,\Gamma_k(t_k,\cdot)),(\Gamma_1(s_1,\cdot),\ldots,\Gamma_k(s_k,\cdot))\rangle_R \\ = \Gamma_1(t_1,s_1)+\ldots+\Gamma_k(t_k,s_k).\end{aligned}$$

and

$$\begin{aligned}\langle(f_1,\ldots,f_k),(\Gamma_1(t_1,\cdot),\ldots,\Gamma_k(t_k,\cdot))\rangle_R \\ = (f_1(t_1),\ldots,f_2(t_k))\end{aligned}$$

for all $(f_1, \ldots, f_k) \in R$.

Theorem 1: (*Generalized Schilder's Theorem*) The function $I : \Omega^k \to \mathbb{R} \cup \{\infty\}$,

$$I(\omega) = \begin{cases} \frac{1}{2}\|\omega\|_R^2, & \text{if } \omega \in R, \\ \infty, & \text{otherwise,}\end{cases}$$

is a good rate function for the centered Gaussian measure P, and P satisfies the large deviation principle:

$$\begin{aligned}&\text{for } F \text{ closed in } \Omega^k: \\ &\quad \limsup_{k\to\infty}\frac{1}{k}\log \mathrm{P}\left(\frac{Z}{\sqrt{k}} \in F\right) \quad \le -\inf_{\omega\in F} I(\omega); \\ &\text{for } G \text{ open in } \Omega^k: \\ &\quad \liminf_{k\to\infty}\frac{1}{k}\log \mathrm{P}\left(\frac{Z}{\sqrt{k}} \in G\right) \quad \ge -\inf_{\omega\in G} I(\omega).\end{aligned}$$

One often distinguishes between two kinds of large deviation asymptotic regimes for queueing systems. In "large buffer" asymptotics one studies the probability that a very large buffer level is exceeded: $\mathrm{P}(Q > x)$, $x \to \infty$. In "many sources" asymptotics, one lets the input be a superposition of n independent, identically distributed streams, and multiplies the server capacity c and the considered buffer level x by n as well. Our asymptotics are of the latter type, as seen in the next lemma.

Lemma 2: (*Many source LDP*) Consider a GPS queue defined by (2) and (3) but replace the input process vector by a superposition of n i.i.d. copies of $(A^{\{1\}}, \ldots, A^{\{k\}})$ and server capacity c by nc. Denote the corresponding class-wise queue length process vectors by $(Q^{(1,n)}, \ldots, Q^{(k,n)})$. Then, for any n,

$$\mathrm{P}\left(Q_0^{(1,n)} \ge nx_1, \ldots, Q_0^{(k,n)} \ge nx_k\right) = \mathrm{P}\left(\frac{Z}{\sqrt{n}} \in B\right),$$

where

$$B = \left\{Q_0^{\{1\}} \ge x_1, \ldots, Q_0^{\{k\}} \ge x_k\right\}.$$

Proof: Note first that in our Gaussian case we can equivalently use $\sqrt{n}Z_t + nt(m_1, \ldots, m_k)$ as the input process. The result is then seen by writing the implicit queueing process definitions (2) and (3), and dividing by n. ■

The essential problem is to find a path ω that minimizes $I(\omega)$ in a given set. We call it the *most probable path* in that set. In most cases of interest, the most probable path is unique. Identifying most probable paths is interesting with its own rights — it is like "seeing what really happens" when the rare event occurs. For ordinary queues, this has mainly heuristic value, but we shall see that finding these paths has an essential role in choosing a good approximation in the case of GPS queues. Further, it is shown in [15], [16] by examples of ordinary queues that the large deviation estimate is often a reasonable approximation for the whole queue length distribution, not only for the tail behavior.

Definition 1: Let $E \subset \Omega^k$ be a set. A *most probable path* $\omega_E \in \overline{E}$ (closure of E) satisfies $I(\omega_E) \le I(\omega)$, $\forall \omega \in E$, and the corresponding *basic approximation* is given by $\mathrm{P}(E) \approx e^{-I(\omega_E)}$.

Remark: The most probable paths correspond to the most probable paths of centered Gaussian processes! The deterministic trends $m_i t$ must be added afterwards.

III. Approximations for the queue length distributions

A. Aggregated queue

In this paper, we consider sets of the form $E = \{Q_0^{\{i\}} \ge x\}$, that is, threshold exceedance events. The simplest problem is to find the most probable class-wise paths for which the total queue exceeds a value x. Since a superposition of independent Gaussian processes is still a Gaussian process, this problem can be reduced to a FIFO queue and the class-wise paths are the only extra information we are seeking here compared to [15], [16].

Proposition 2: The most probable path vectors in the set $\{Q_0^{\{1,\ldots,k\}} \ge x\}$ have the form

$$f_x^*(\cdot) = -\frac{x + (c - \sum_{i=1}^k m_i)(-t_x)}{\sum_{i=1}^k v_i(t_x)}(\Gamma_1(t_x,\cdot), \ldots, \Gamma_k(t_x,\cdot)),$$

where $t_x < 0$ is the value of t which minimizes

$$h(t) = \frac{(x + (c - \sum_{i=1}^k m_i)(-t))^2}{\sum_{i=1}^k v_i(t)}. \tag{5}$$

If v_i is differentiable, then the minimizing problem is equivalent to finding solutions of

$$\frac{\sum_{i=1}^k v_i(t)}{\sum_{i=1}^k v_i'(t)} = \frac{1}{2}\left(t - \frac{x}{c - \sum_{i=1}^k m_i}\right).$$

Proof: Note that

$$\begin{aligned}\left\{Q_0^{\{1,\ldots,k\}} \ge x\right\} &= \bigcup_{t\le 0}\left\{A^{\{1,\ldots,k\}}(t,0) - c(0-t) \ge x\right\} \\ &= \bigcup_{t\le 0}\left\{Z^{\{1,\ldots,k\}}(t,0) + (c-m)t \ge x\right\},\end{aligned}$$

and, by the reproducing kernel property,

$$\begin{aligned} &f \in \left\{Z^{\{1,\dots,k\}}(t,0) \geq x-(c-m)t\right\} \cap R && \Leftrightarrow \\ &f \in R,\ f_1(t)+\cdots+f_k(t) \leq -x+(c-m)t && \Leftrightarrow \\ &\langle f,(\Gamma_1(t,\cdot),\dots,\Gamma_k(t,\cdot))\rangle_R \leq -x+(c-m)t. \end{aligned}$$

Thus, the problem reduces to minimizing the Hilbert norm when the inner product with a fixed element is given, and the solution is a proper multiple of that element. It remains to minimize the norm of

$$-\frac{x-(c-m)t}{\sum v_i(t)}(\Gamma_1(t,\cdot),\dots,\Gamma_k(t,\cdot))$$

with respect to t. ■

A lower bound for the aggregated queue distribution can be derived using the 1-dimensional normal distribution:

$$\begin{aligned} &\mathrm{P}\left(Q_0^{\{1,\dots,k\}} \geq x\right) \\ &\geq \mathrm{P}\left(A_{-t_x}^{\{1,\dots,k\}} \geq x-(c-\sum_{i=1}^k m_i)t_x\right) \\ &= \overline{\Phi}\left(\frac{x-(c-\sum_{i=1}^k m_i)t_x}{\sqrt{\sum_{i=1}^k v_i(t_x)}}\right), \end{aligned} \tag{6}$$

where $t_x < 0$ is again the value of t which minimizes (5).

B. Class 1 queue

Consider next the queueing process of class 1. Since classes have a symmetric role in a GPS queue, this restricts nothing.

Let f_x^* be the most probable path vector in the set $\{Q_0^{\{1,\dots,k\}} \geq x\}$ given by Proposition 2. There are two different ways how the class 1 queue can build up.

Case 1:

$$Q_0^{\{2,\dots,k\}}(f_x^*) = 0$$

(implying that $Q_0^{\{1\}}(f_x^*) \geq x$). Then the combined queue equals the queue of class 1, and it follows that this path is also the most probable one to achieve level x in the class 1 queue alone. (This is equivalent to the "Empty buffer approximation" introduced by Berger and Whitt in [23], [24].)

Case 2:

$$Q_0^{\{1\}}(f_x^*) < x.$$

The situation is more difficult, and the most probable paths remain unknown, except for the 2 class Brownian case discussed in the next example.

Note that no more than one class can belong to Case 1 at a time.

Example 1: If a GPS queue has two independent Brownian motions as input, i.e., $v_i(t) = C_i t$, with C_i constant, $i = 1, 2$, and the parameters are such that the system belongs to Case 2, then it can be shown that the most probable paths in the set $\{Q_0^{\{1\}} \geq x\}$ satisfy the following:

- Class 1: $f^{\{1\}}(t_x, 0) = x + (\mu_1 c - m_1)(-t_x)$
- Class 2: $\frac{d}{dt} f_t^{\{2\}} = \mu_2 c$ if $t \in (t_x, 0)$, i.e., the input rate of class 2 is exactly its guaranteed rate.
- $t_x < 0$ is the optimal time under the previous conditions.

Motivated by the Brownian example, it might be reasonable to use an approximation where the other classes do not queue at all, albeit they use their reserved capacity fully (*Full Link Approximation*). Unfortunately, this is usually uncomputable too. Thus we suggest to relax the exact rate demand and replace it by a condition for the total amount of traffic offered (*Rough Full Link Approximation*).

Definition 2: The *Rough Full Link Approximation* (RFLA) for a GPS queue with two input classes. During $[t, 0]$, $t < 0$,

- class 1 offers in total the amount $x + \mu_1 c|t|$ of traffic: $A^{\{1\}}(t, 0) = x + \mu_1 c|t|$,
- class 2 offers in total the amount $\mu_2 c|t|$ of traffic: $A^{\{2\}}(t, 0) = \mu_2 c|t|$,
- $[t_x, 0]$ is the most probable interval under the above conditions.

In principle, the idea above could be extended to a larger number of classes, but the details would be much more complicated and, moreover, the heuristic probability estimates would be less reliable.

It is an easy Hilbert space exercise to determine the most probable paths in RFLA.

Proposition 3: The most probable paths satisfying the RFLA conditions are of the form

$$f_x^{\{1\}}(\cdot) = \begin{pmatrix} \frac{-x+(\mu_1 c - m_1)t_x}{v_1(t_x)}\,\Gamma_1(t_x,\cdot) \\ \frac{(\mu_2 c - m_2)t_x}{v_2(t_x)}\,\Gamma_2(t_x,\cdot) \end{pmatrix}$$

where $t_x < 0$ is the value of t which minimizes the expression

$$g_1(t) = \frac{(x-(\mu_1 c - m_1)t)^2}{v_2(t)} + \frac{(\mu_2 c - m_2)^2 t^2}{v_2(t)}$$

Proof: If $(f_1, f_2) \in R$ and they satisfy the RFLA conditions, then by the reproducing kernel property

$$\begin{aligned} \langle f_1, \Gamma_1(t,\cdot)\rangle_{R_1} &= f_1(t) = (\mu_1 c - m_1)\,t - x \\ \langle f_2, \Gamma(t,\cdot)\rangle_{R_2} &= f_2(t) = (\mu_2 c - m_2)\,t. \end{aligned}$$

Thus, the most probable path must be of the form

$$\left(\frac{f_1(t)}{v_1(t)}\Gamma_1(t,\cdot), \frac{f_2(t)}{v_2(t)}\Gamma_2(t,\cdot)\right).$$

Minimizing the R-norm with respect to t concludes the proof. ■

Now, we have the following approximations for the class 1 queue length distribution:

Case 1: The basic approximation

$$\mathrm{P}\left(Q^{\{1\}} \geq x\right) \approx \mathrm{P}\left(Q^{\{1,2\}} \geq x\right) \approx \exp\left(-\frac{1}{2}h(t_x)\right)$$

and the corresponding lower bound (6).

Case 2: The basic approximation

$$\mathrm{P}\left(Q^{\{1\}} \geq x\right) \approx \mathrm{P}(\mathrm{RFLA}(x)) \approx \exp\left(-\frac{1}{2}g_1(t_x)\right).$$

and a lower bound for the RFLA

$$\begin{aligned} &\mathrm{P}(\mathrm{RFLA}(x)) \geq \\ &\quad \overline{\Phi}\left(\frac{x-(\mu_1 - m_1)t_x}{\sqrt{v_1(t_x)}}\right)\overline{\Phi}\left(-\frac{(\mu_2 c - m_2)t_x}{\sqrt{v_2(t_x)}}\right). \end{aligned}$$

Simulations of many cases indicate that the basic approximations for the tail distributions may in fact be general upper bounds (see Section IV). Unfortunately, no proof of this is known, not even in the FIFO case.

C. Improving approximations by rescaling

If a system is such that the buffer is empty most of the time (which is typical for most telecommunication systems), the basic approximations overestimate quite much. Moreover, if we compare approximations with empirical distributions from discrete simulations, there is an extra gap because of the discretization. A natural approach would be to rescale the basic approximations so that the non-emptiness probabilities $\mathrm{P}\left(Q^{\{i\}} > 0\right)$ are estimated as well as possible.

In [15], [16], upper and lower bounds of $\mathrm{P}(Q > 0)$ were derived in the case of ordinary queues. It was found that a good approximation for non-priority queues is $\mathrm{P}(Q > 0) \approx 2\mathrm{P}(A_{\Delta t} > c\Delta t)$, where Δt can be interpreted as a limit, a discretization step or the smallest time scale to approve a Gaussian model. The idea of the approximation is that when $A_{\Delta t} > c\Delta t$ then the buffer is non-empty and going up, and since the buffer is roughly as often non-empty and going down, we multiply the probability by 2.

Unfortunately, the corresponding approximations are not known for GPS systems. In order to maintain the (heuristic) upper bound characteristic of the basic approximations, we consider the following worst-case scenario. Class i queue cannot be larger than the total queue. On the other hand, class i queue is always smaller than the queue of a single class system with input $A^{\{i\}}$ and server rate $\mu_i c$. Applying the non-emptiness approximation into these two non-priority queues gives

$$\begin{aligned}\mathrm{P}\left(Q^{\{i\}} > 0\right) &\lesssim \\ & 2\min\left\{\overline{\Phi}\left(\frac{(c-\sum m_j)\,\Delta t}{\sqrt{\sum v_j(\Delta t)}}\right), \overline{\Phi}\left(\frac{(\mu_i c - m_j)\,\Delta t}{\sqrt{v_i(\Delta t)}}\right)\right\} \\ &= p_i(\Delta t),\end{aligned}$$

where $\overline{\Phi}$ denotes the complementary normal distribution.

Thus the rescaled basic approximations having approximately correct non-emptiness probabilities can be written as follows.

Case 1:

$$\mathrm{P}\left(Q^{\{i\}} \geq x\right) \approx p_i(\Delta t)\exp\left(-\frac{1}{2}(h(t_x) - \lim_{x\to 0} h(t_x))\right)$$

Case 2:

$$\mathrm{P}\left(Q^{\{i\}} \geq x\right) \approx p_i(\Delta t)\exp\left(-\frac{1}{2}(g_i(t_x) - \lim_{x\to 0} g_i(t_x))\right).$$

IV. Example of a GPS system analysis

In this section, we show how the approach introduced above can be used to analyze a GPS system. In real-life applications, one should be able to measure mean rates, validate the Gaussian assumption and fit proper Gaussian models to each class. How to perform the two last steps is far from trivial. A preliminary solution is given in [2]. Here, we start from Gaussian traffic and do not bother about the practical issues of Gaussian modeling.

In addition to demonstrating the applicability of the performance estimates, we want to stress the following important aspects of GPS schedulers with Gaussian input that came out in our study:

- The system is very sensitive; is some regions small changes of parameters may change performance a lot.
- Mean rate based weight assignments do not usually give desired results. It is important to take into account the variance structure carefully.
- Fairness is difficult to obtain/define; one can not assume similar queueing behavior – not even qualitatively – for processes which have different types of variances.

A. Input traffic

Let us consider a GPS node serving two independent Gaussian flows which are characterized by mean rates

$$m_1 = m_2 = 2$$

and variance functions

$$v_1(t) = 4(t - 1 + e^{-t}), \quad v_2(t) = t^{\frac{3}{2}}.$$

Assume unlimited buffers and the total server rate $c = 5$. In this example, class 1 traffic is short-range dependent and class 2 traffic is long-range dependent.

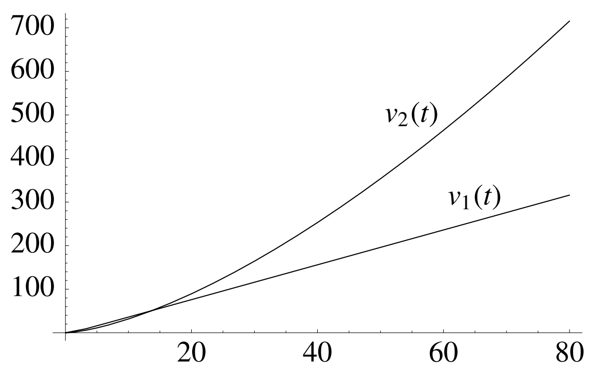

Fig. 1. Variance functions.

The variance functions are shown in Figure 1. The order of variances depends on the value of t: v_1 is larger than v_2 on $[0.3, 13.8]$, otherwise v_2 majorizes. We should expect that the very smallest and the largest queues are more likely built by class 2 traffic, since the larger the variance the "easier" it is to create bursts. In the intermediate region we have a kind of mixed situation. Naturally, the GPS weights affect on the proportion of the classes in the total queue, but qualitatively the above reasoning should hold for all weight pairs. This can be seen also in the most probable path approach and simulations shown in the next subsections.

B. Most probable input rates and storage paths

As already mentioned, the most probable paths leading to certain total queue level do not depend on the GPS weights at

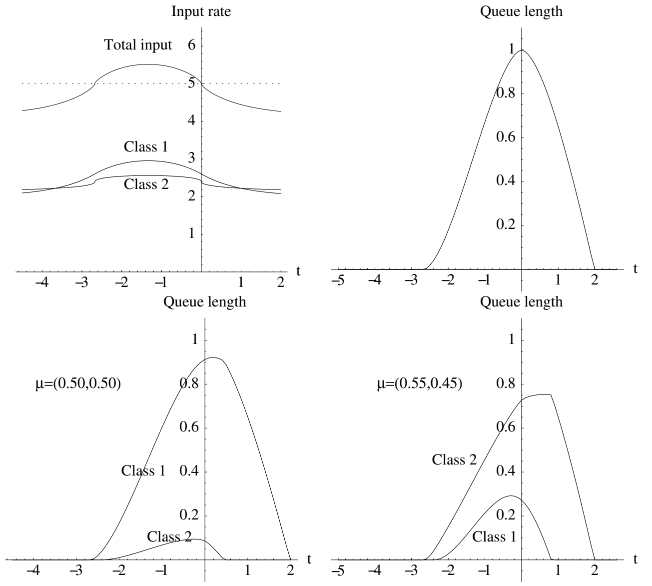

Fig. 2. $\{Q_0^{\{1,2\}} \geq 1\}$. The most probable input rates, the corresponding total and class-wise queues.

all. The most probable input rates in $\{Q_0^{\{1,2\}} \geq 1\}$ and the aggregated queue are plotted in Figure 2. On the other hand, the weights have a strong effect on the queueing behavior of different classes. The most probable storage paths corresponding to two slightly different weight vectors are also shown in Figure 2. We see that the most probable storage paths have changed quite a lot: class 1 dominates the aggregated queue if $\mu = (0.50, 0.50)$, whereas class 2 queue is the larger one if $\mu = (0.55, 0.45)$.

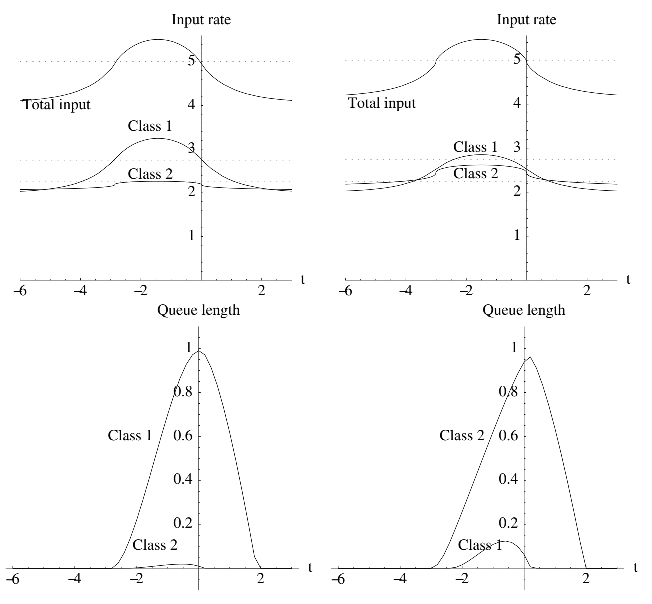

Fig. 3. The most probable input rates and corresponding queueing processes in the RFLA for $x = 1$ and $\mu = (0.55, 0.45)$. The dashed lines correspond to the service guarantees.

Figure 2 shows that both classes belong to Case 2 when $x = 1$ and $\mu = (0.55, 0.45)$ or $\mu = (0.50, 0.50)$. Thus, in order to approximate class-wise tail distributions we have to use RFLA. The most probable input rates and the corresponding queues in the RFLA's for the both classes are shown in Figure 3. The approximate character of the RFLA conditions is clearly seen here; the target queue does not usually reach quite the buffer level x, and, moreover, the other class has often a small queue left at $t = 0$.

C. Comparison to simulations

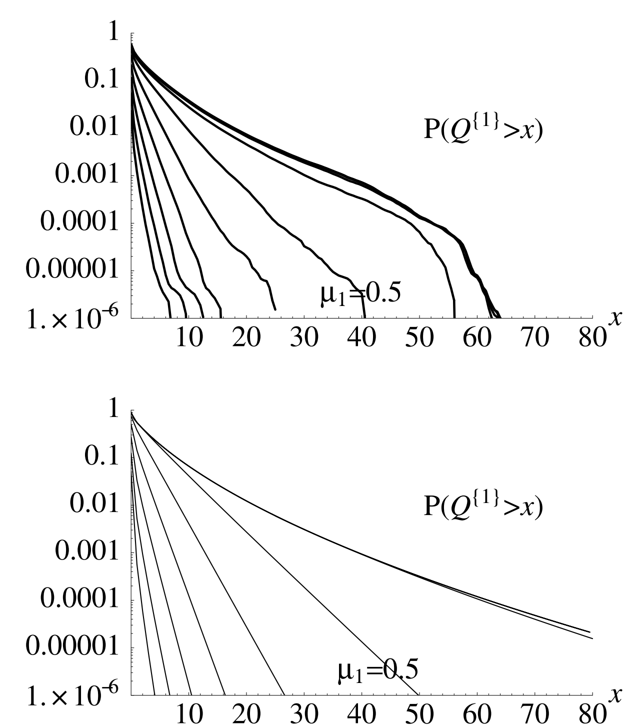

Fig. 4. Simulation results for the class 1 queue and the corresponding basic approximations: $\mu_1 = 0.0, 0.1, 0.2, \ldots, 1.0$. Simulation lengths 2^{23} steps at resolution 0.125.

In order to check how well the approximations work we compare them with simulations. The Gaussian traces were generated using an extension of the conditionalized random midpoint displacement algorithm (RMD$_{mn}$, see [25]).

First we study the qualitative behavior of the basic approximations. Let us consider the weights $\mu_1 = 0, 0.1, \ldots, 0.9, 1.0$. Note that the pure priorities are included. By examining the most probable storage paths, we determine which formulae to apply. The approximations are calculated for buffer lengths $x \in [0.5, 80]$. Class 1: RFLA if $\mu_1 \leq 0.5$ or $\mu_1 = 0.6$ and $x > 25$, otherwise the Case 1 approximations. Class 2: RFLA if $\mu_2 \leq 0.4$ or $\mu_2 = 0.5$ and $x < 32$, otherwise the Case 1 approximations. The simulation results and the corresponding basic approximations are shows in Figures 4 and 5. The approximations perform surprisingly well. Only when $\mu_2 = 0.5$ the empirical measure and the approximation have a bit different flavor, since the RFLA based estimate is indistinguishable from the total queue estimate.

In Figure 6, we have plotted all the approximations for two different setups. The empirical values stay well in between our

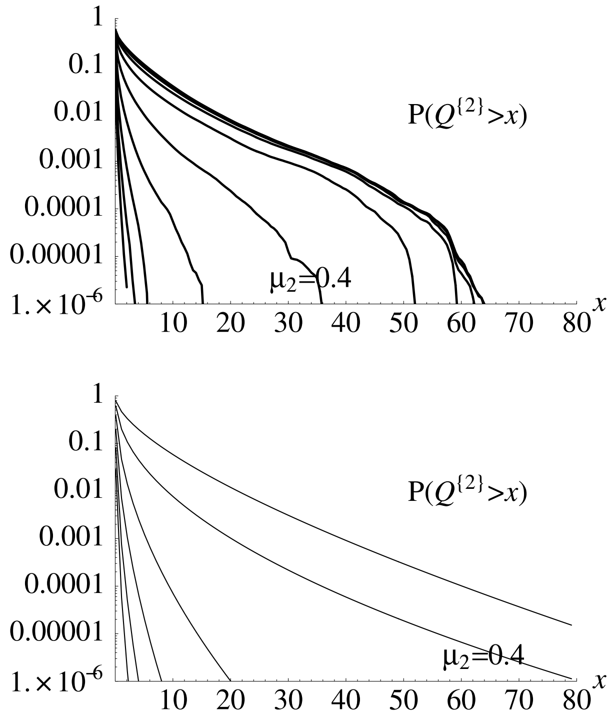

Fig. 5. Simulation results for the class 2 queue and the corresponding basic approximations: $\mu_2 = 0.0, 0.1, 0.2, \ldots, 1.0$. Simulation lengths were 2^{23} steps at the resolution 0.125.

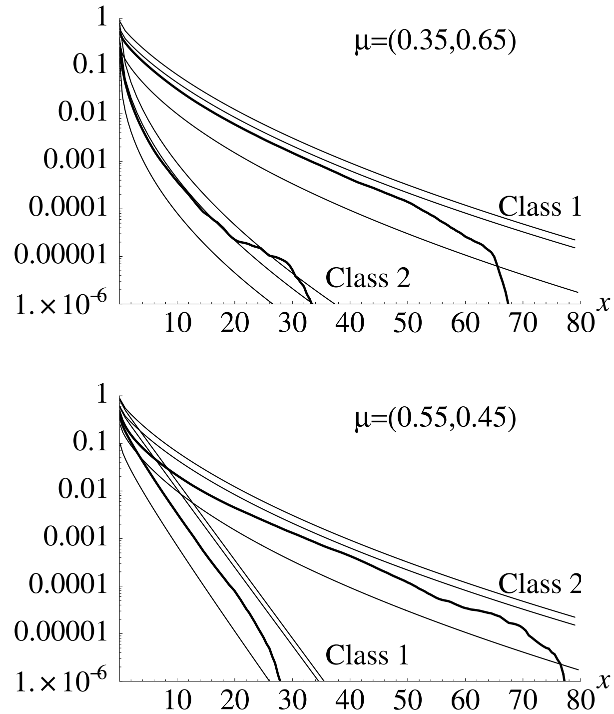

Fig. 6. Approximations of the queue length distributions. The thick curves are the empirical distributions from the simulations. The approximations from top down: the basic approximation, the rescaled version of it and the corresponding (heuristic) lower bound. The simulation length was 2^{24} steps at the resolution 0.25.

estimates until the typical drops at the ends, caused by finiteness of the simulations. The accuracy is not extremely good in the sense of closeness between upper and lower estimates, but certainly sufficient for finding essential differences in the queue length distributions.

When considering the behavior of Gaussian GPS queues, the immediate lesson we get from Figures 4 and 5 is that the system can be very sensitive with respect to the weights. If the parameters are such that the class does not belong to Case 1, small changes may alter its behavior radically. Whereas classes in Case 1 may need a remarkable increase of the service guarantee in order to get better performance.

Another property is that the variances have a great effect on the performance. In our examples, Class 2 traffic suffers clearly from the long range dependence. E.g., the mean rate based fair share of the server, i.e., $\mu_1 = \mu_2 = 0.5$, does not result in fair queueing as seen in the figures.

How to determine weights fairly? Because of the different correlation ranges, the queue length distributions are always qualitatively different. One solution could be to determine the weights so that the total queue of size x is built equally by both classes. In Figure 7, it is shown what is the most probable way to build a total queue of size $x = 30$. If we have equal weights, the queue consists almost only of Class 2 traffic, whereas weight pair $\mu = (0.45, 0.55)$ results in equal queue lengths. How does this affect on the queue length distributions? The only thing that can be said before plotting the approximations is that the intersection point should be somewhere with $x < 30$. In this case, the basic estimates suggest that the inter-

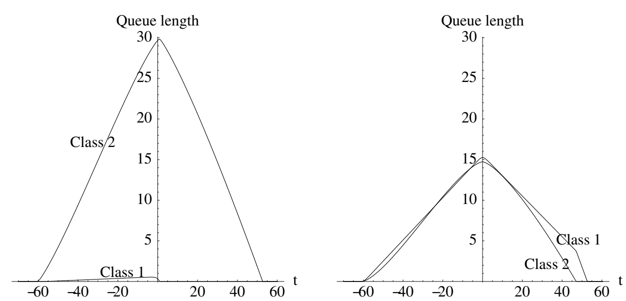

Fig. 7. The most probable storage paths in $\{Q_0^{\{1,2\}} \geq 30\}$ for $\mu = (0.5, 0.5)$ and $\mu = (0.45, 0.55)$.

section happens around $x = 25$, which is in accordance with the results from the simulation (see Figure 8).

V. Concluding remarks

We have introduced a most probable path approach which can be used to analyze GPS schedulers with Gaussian input. Although the method is computationally quite simple, it gives crucial information about the behavior of the queueing systems. Especially, relating queue length distributions with GPS weights can be done easily with this machinery.

The example we analyzed shows that a GPS system is very sensitive and a mild modification of the parameters may produce large changes in the queueing behavior. Moreover, there are many limitations how fairly the classes can be served. For

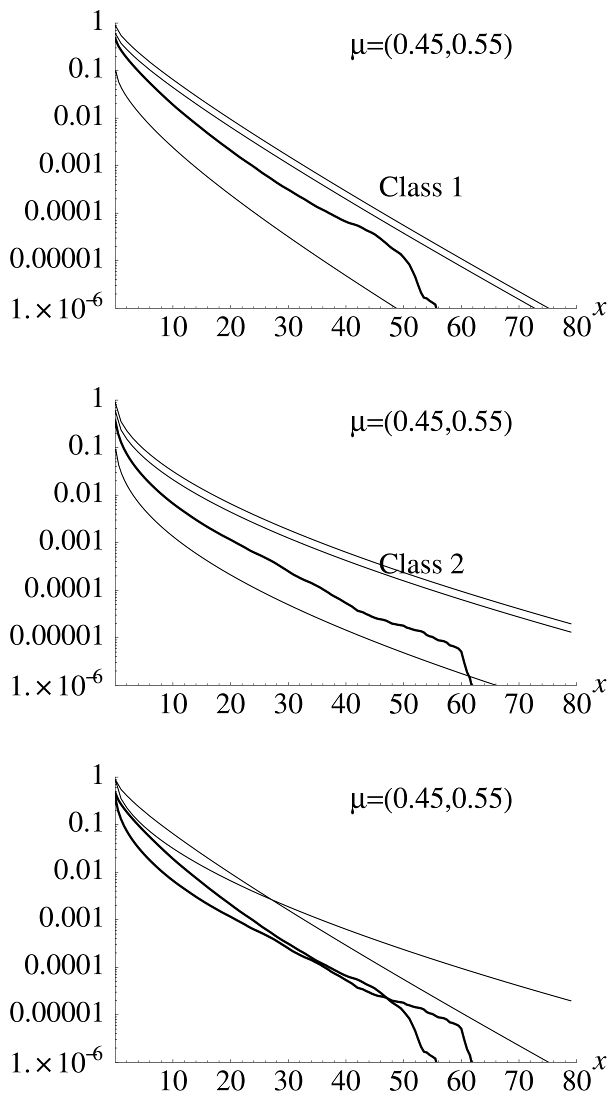

Fig. 8. Approximations of the queue length distributions. The thick curves are the empirical distributions from the simulation. The approximations from top down: the basic approximation, rescaled version of it and the corresponding (heuristic) lower bound. In the lowest figure, only the simulations and the basic approximations are shown. The simulation length 2^{24} was steps at the resolution 0.25.

example, giving all classes an identical minimum rate guarantee may effect that the waiting time distributions are totally unfair.

This small study strengthens criticism about using methods that require fine-tuning of parameters for resource allocations in Differentiated Services Networks. It will be very hard to build complex systems with several service classes and to be still able to give QoS guarantees which can be confirmed. Even if the traffic characteristics were known exactly, the specification of the service parameters can be very difficult. However, in real networks we can only hope that our models are somewhat near to real traffic, and as shown also in this paper, the resource allocation systems may be quite unstable. Thus, simple self-adjusting control schemes, like Kalevi Kilkki's SIMA [26], should be taken into consideration when designing future networks.

References

[1] R.G. Addie, "On the weak convergence of long range dependent traffic processes," *J. Statistical Planning and Inference*, vol. 80, pp. 155–171, 1999.

[2] J. Kilpi and I. Norros, "Testing the Gaussian character of access network traffic," COST279TD(01)03, VTT Information Technology, 2001, `http://www.vtt.fi/tte/projects/cost279/`.

[3] K. Kilkki, *Differentiated services for the Internet*, Macmillan Technical Publishing, Indiapolis, USA, 1999.

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

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

[6] Z.-L. Zhang, Z. Liu, and D. Towsley, "Closed-form deterministic performance bounds for the generalized processor sharing scheduling discipline," *Journal of Combinatorial Optimization*, vol. 1, no. 4, 1998, Special Issue on Scheduling.

[7] G. de Veciana and Kesidis G., "Bandwidth allocation for multiple qualities of services using generalized processor sharing," *IEEE Trans. Inform. Theory*, vol. 42, no. 1, pp. 268–272, 1996.

[8] Z.-L. Zhang, "Large deviations and the generalized processor sharing scheduling for a two-queue system," *Queueing Systems*, vol. 26, pp. 229–254, 1997.

[9] Z.-L. Zhang, "Large deviations and the generalized processor sharing scheduling for a multiple-queue system," *Queueing Systems*, vol. 28, pp. 349–376, 1998.

[10] D. Bertsimas, I.C. Paschalidis, and J.N. Tsitsiklis, "Large deviations analysis of generalized processor sharing policy," *Queueing Systems*, vol. 32, pp. 319–349, 1999.

[11] L. Massoulié, "Large deviations estimates for polling and weighted fair queueing service systems," *Advances in Performance Analysis*, vol. 2, no. 2, pp. 103–127, 1999.

[12] J.-Y. Le Boudec and P. Thiran, *A theory of deterministic queueing systems for the Internet*, Number 2050 in LNCS. Springer, 2001.

[13] I. Norros, "A storage model with self-similar input," *Queueing Systems*, vol. 16, pp. 387–396, 1994.

[14] I. Norros, "Busy periods of fractional Brownian storage: a large deviations approach," *Advances in Performance Analysis*, vol. 2, no. 1, pp. 1–20, 1999.

[15] R.G. Addie, P. Mannersalo, and I. Norros, "Performance formulae for queues with Gaussian input," in *Proceedings of ITC 16*, P. Key and D. Smith, Eds. 1999, vol. 3b, pp. 1169–1178, Elsevier.

[16] R.G. Addie, P. Mannersalo, and I. Norros, "Most probable paths and performance formulae for buffers with Gaussian input traffic," *European Transactions on Telecommunication*, 2002, To appear.

[17] P. Mannersalo and I. Norros, "Gaussian priority queues," in *Proceedings of ITC 17*. Elsevier, 2001, pp. 991–1002.

[18] I. Norros, "A storage model with self-similar input," *Queueing Systems*, vol. 16, pp. 387–396, 1994.

[19] R.R. Bahadur and S.L. Zabell, "Large deviations of the sample mean in general vector spaces," *Ann. Prob.*, vol. 7, no. 4, pp. 587–621, 1979.

[20] R. Azencott, *Ecole d'Eté de Probabilités de Saint-Flour VIII-1978*, chapter Grandes deviations et applications, pp. 1–176, Number 774 in Lecture notes in Mathematics. Springer, Berlin, 1980.

[21] J.-D. Deuschel and D.W. Stroock, *Large Deviations*, Academic Press, 1989.

[22] R.J. Adler, *An Introduction to Continuity, Extrema, and Related Topics for General Gaussian Processes*, IMS Lecture Notes-Monographs Series. Institute of Mathematical Statistics, 1990.

[23] A.W. Berger and W. Whitt, "Effective bandwidths with priorities," *IEEE/ACM Transaction on Networking*, vol. 6, no. 4, pp. 447–460, 1998.

[24] A.W. Berger and W. Whitt, "Extending the effective bandwidth concept to networks with priority classes," *IEEE Communications Magazine*, pp. 78–83, August 1998.

[25] I. Norros, P. Mannersalo, and J.L. Wang, "Simulation of fractional Brownian motion with conditionalized random midpoint displacement," *Advances in Performance Analysis*, vol. 2, no. 1, 1999.

[26] K. Kilkki, "Simple Integrated Media Access (SIMA)," Internet-Draft, see `http://www.nokia-nrc.com/sima/`.