

Most Probable Path Techniques for Gaussian Queueing Systems

Ilkka Norros

VTT Information Technology, P.O. Box 1202, 02044 VTT, Finland

Abstract. This paper is a review of an approach to queueing systems where the cumulative input is modelled by a general Gaussian process with stationary increments. The examples include priority and Generalized Processor Sharing systems, and a system where service capacity is allocated according to predicted future demand. The basic technical idea is to identify the most probable path in the threshold exceedance event, or a heuristic approximation of it, and then use probability estimates based on this path. The method is particularly useful for long-range dependent traffic and complicated traffic mixes, which are difficult to handle with traditional queueing theory.

1 Introduction

This paper is a review of an approach to queueing systems with Gaussian input. The motivation to study such systems is twofold. On one hand, complicated dependence structures are easiest to study first in a Gaussian framework, where the dependence is reduced to correlation. This is also the historical origin of this work — it started with queues with fractional Brownian motion (fBm) as input [19], which is the simplest process that has the self-similarity property, first observed in the famous Bellcore measurements [11]. On the other hand, it could be expected that, thanks to the Central Limit Theorem, traffic in high capacity systems would be rather well modelled with Gaussian processes [1]. Empirical studies indicate, however, that a good fit to Gaussian distribution may require very high traffic aggregation levels. The Gaussian approach can be useful in making rough performance estimates for Differentiated Services in Internet, because one works there with large traffic aggregates.

Our interest in most probable paths started by applying the generalized Schilder's theorem to the fBm queue [20]. The approach was extended to ordinary queues with general Gaussian input in [2,3], and further to priority queues in [14]. In [13] and [15], we applied a similar machinery to Generalized Processor Sharing (GPS) schedulers and presented a somewhat improved version of the priority case. Most of this research was done within the COST Actions 257 and 279. A summary on Gaussian traffic modelling, linked to the technical documents, can be found in the hypertext Final Report of the action [26].

The paper is structured as follows. We start with discussing the definitions of Gaussian queueing systems in Section 2. This involves some technical details caused by the unavoidable presence of negative traffic in Gaussian modelling.

Section 3 presents the main ideas of our approach. A central role is played by the most probable paths along which queue size thresholds are exceeded. The rest is devoted to two cases, where the most probable paths obtain particularly interesting shapes. Section 4 shows how the most probable path can experience a kind of “phase transition” between short and long busy periods. Section 5 studies a simple model of dynamical capacity allocation. This is a new type of application, first time presented here.

2 Definition of Gaussian Queueing Systems

2.1 Gaussian Models of Traffic

Our basic traffic model is a continuous Gaussian process $A = (A_t)_{t \in \mathfrak{R}}$ with stationary increments. For $s < t$, $A_t - A_s$ presents the amount of traffic in time interval $(s, t]$, and we set $A_0 \equiv 0$. A process is called Gaussian, if all its finite-dimensional distributions are multivariate Gaussian. The property of stationary increments means that for any $t_0 \in \mathfrak{R}$, the processes A and $(A_{t+t_0} - A_{t_0})_{t \in \mathfrak{R}}$ have the same finite-dimensional distributions.

We denote $A(s, t) = A_t - A_s$, and use similar notation for other processes as well.

The use of Gaussian models for big traffic aggregates can be justified by the Central Limit Theorem. However, even the question about the Gaussian character of some traffic cannot be raised without specifying the relevant timescale, say δ . There should be a large number of individual sources contributing to the traffic in every time interval of size δ . Moreover, if the marginal distribution of the contribution of an individual source in those intervals has very high variability, the application of CLT may still be problematic. In our study on Internet users over ISDN [10], it was found that a few Mbit/s of such traffic had good fit with Gaussian distribution when the time resolution δ was coarser than 100 ms. Note that this traffic was exceptionally well-behaving, because the users were restricted to the ISDN access speed.

A non-pleasant special feature of Gaussian models is that there is always a positive probability of negative input. Such input does not correspond to anything real, and its existence destroys some classical arguments of queueing theory. In a Gaussian framework, the non-problematic definitions of queueing theory must be replaced by analogously defined functionals of a Gaussian process. Moreover, we don't have much hope to obtain other kinds of rigorous general results on the distributions of these functionals than inequalities and limit theorems. At our present “state-of-art”, we must often be satisfied with heuristic approximations.

Despite these reservations, Gaussian models are tempting because of their many nice features:

- a Gaussian process with stationary increments is completely characterized by its mean $m = \mathbb{E}\{A_1\}$ and cumulative variance function $v(t) = \text{Var}(A_t)$; indeed, we can write

$$A_t = mt + Z_t,$$

where Z is a centered (mean zero) process, and the covariance function of A (and Z) can be written as

$$\text{Cov}(A_s, A_t) = \text{Cov}(Z_s, Z_t) = \frac{1}{2}(v(s) + v(t) - v(s - t));$$

- a superposition of independent Gaussian traffic streams is Gaussian;
- multiclass traffic consisting of Gaussian traffic classes, such that their joint distribution is Gaussian also, can be studied within the same framework;
- unlike most other traffic models, A_t has an explicitly known (Gaussian) distribution for any t ;
- the quantities m and $v(t)$ can be rather well estimated from measurement data;
- long-range dependence does not provide any extra difficulty.

In the multiclass case, let the input traffic consist of k classes, and denote the cumulative arrival process of class $j \in \{1, \dots, k\}$ by $(A_t^{\{j\}})_{t \in \mathfrak{R}}$. We also denote $A^{\{j\}}(s, t) \doteq A_t^{\{j\}} - A_s^{\{j\}}$. For the superposition of a set of traffic classes $J \subseteq \{1, \dots, k\}$ we write

$$A_t^J \doteq \sum_{j \in J} A_t^{\{j\}}$$

and use similar superscript notation also for other quantities defined later. We assume that the processes $A^{\{j\}}$ are independent, continuous Gaussian processes with stationary increments and denote

$$A_t^{\{j\}} = m_j t + Z_t^{\{j\}}, \quad m = \sum_{i=1}^k m_i, \quad \text{Var}(Z_t^{\{j\}}) = v_j(t), \quad (1)$$

$$\Gamma_j(s, t) = \text{Cov}(Z_s^{\{j\}}, Z_t^{\{j\}}),$$

where the $Z^{\{j\}}$'s are centered (zero-mean) processes. To exclude certain degenerate cases, we assume that

$$\exists \alpha \in (0, 2) : \lim_{t \rightarrow \infty} \frac{v_i(t)}{t^\alpha} = 0, \quad i \in \{1, \dots, k\}. \quad (2)$$

Finally, let us specify the mathematical framework completely. Define a path space Ω_1 as

$$\Omega_1 = \left\{ \omega : \omega \text{ is continuous } \mathfrak{R} \rightarrow \mathfrak{R}, \omega(0) = 0, \lim_{t \rightarrow \pm\infty} \frac{\omega(t)}{1 + |t|} = 0 \right\}.$$

(The relation $\lim_{t \rightarrow \infty} Z_t^{\{i\}}/t = 0$ a.s., is a consequence of (2) — see [3].) Equipped with the norm

$$\|\omega\|_{\Omega_1} = \sup \left\{ \frac{\omega(t)}{1 + |t|} : t \in \mathfrak{R} \right\},$$

Ω_1 is a separable Banach space. We choose $\Omega = \Omega_1^k$ as our basic probability space by letting P be the unique probability measure on the Borel sets of Ω such that the random variables $Z_t^{\{i\}}(\omega_1, \dots, \omega_k) = \omega_i(t)$ form independent Gaussian processes with covariance functions $\Gamma_i(\cdot, \cdot)$.

2.2 Definition of Simple Queues

Consider first the case of a simple queue, i.e. $k = 1$, and let the server have a constant capacity c . The storage process (queue length process) is then naturally defined as

$$Q_t = \sup_{s \leq t} (A(s, t) - c(t - s)). \quad (3)$$

The process Q is obviously stationary, and a sufficient stability condition is that $m < c$.

Because only the net input process $A_t - ct$ matters in this definition, it can be extended to the case that the service process is stochastic as well. Indeed, assume that the cumulative service capacity process C_t is a Gaussian process with stationary increments such that the difference $A_t - C_t$ is also Gaussian with stationary increments and a negative mean rate. The queue length process is then

$$Q_t = \sup_{s \leq t} (A(s, t) - C(s, t)), \quad (4)$$

and all results for simple Gaussian queues are applicable. One example of this is given in Section 5.

2.3 Definitions of GPS and Priority Queues

The Generalized Processor Sharing (GPS) service discipline [23] (an idealized version of Weighted Fair Queueing) is a theoretical model which isolates flows and provides service differentiation. Let us consider a GPS queueing system for our k traffic classes, such that the guaranteed service rate for each class i is $\mu_i c$, where $c > m = \sum_i m_i$, $\mu_i > 0$ for each i , and $\sum \mu_i = 1$.

It is not at all obvious how a GPS queue should be defined when negative input is allowed. An elegant definition which results in positive queue length processes even in our case was given by Massoulié [16]. Assume that the amount of potential service for each class i in time interval (s, t) is $\mu_i c T(s, t)$, where $T(s, t) = T_t - T_s$ and T is a non-decreasing stochastic process with $T_0 \equiv 0$. T varies according to the number of backlogged classes. The queue of class i , $Q^{\{i\}}$, and the total queue Q then satisfy

$$\begin{aligned} Q_t^{\{i\}} &= \sup_{s \leq t} (A^{\{i\}}(s, t) - \mu_i c T(s, t)) \\ Q_t &= \sup_{s \leq t} \left(\sum_{i=1}^k A^{\{i\}}(s, t) - c(t - s) \right). \end{aligned} \quad (5)$$

Together with the requirement $Q_t = \sum_{i=1}^k Q_t^{\{i\}}$, the equations (5) uniquely define the $k + 1$ processes $Q^{\{1\}}, \dots, Q^{\{k\}}$ and Q_t [16]. The construction works and yields non-negative queues in the Gaussian case also.

Let us then turn to priority queues. Assume that there are k priority classes, numbered with descending priority. There is no distinction between preemptive and non-preemptive priority, because the model is continuous. Since lower class traffic does not disturb upper class traffic, a simple approach is the following:

define $Q^{\{1\}}$, $Q^{\{1,2\}}$, $Q^{\{1,2,3\}}$ etc. as ordinary queues with service rate c , and then set

$$\begin{aligned} Q^{\{2\}} &= Q^{\{1,2\}} - Q^{\{1\}}, \\ &\dots \\ Q^{\{k\}} &= Q^{\{1,\dots,k\}} - Q^{\{1,\dots,k-1\}}. \end{aligned} \tag{6}$$

Using this definition with Gaussian traffic has the non-desirable effect that it does not yield non-negative queue lengths to other classes than the first one. This has, however, little significance in the cases where Gaussian modeling is adequate, so we prefer using it. (Massoulié's GPS definition does not work, as such at least, with $\mu_2 = 0$, which would correspond to a two-class priority queue. It is shown in [15] how discrete time Gaussian priority queues can be defined in such a way that the individual queues are non-negative and sum up to the total queue, and the continuous time could probably be obtained as a limit when the discretization step goes to zero.)

3 Probability Estimates Based on Most Probable Paths

3.1 The Reproducing Kernel Hilbert Space and Large Deviations of Gaussian Processes

For $i = 1, \dots, k$, the reproducing kernel Hilbert space (RKHS) R_i of the process $Z^{\{i\}}$ is defined as follows (see, e.g., [4]): start with the functions $\Gamma_i(t, \cdot)$, $t \in \mathfrak{R}$, define their inner products as

$$\langle \Gamma_i(s, \cdot), \Gamma_i(t, \cdot) \rangle_{R_i} \doteq \Gamma_i(s, t),$$

extend to a linear space (with pointwise operations), and complete the space with respect to the norm $\|f\|_{R_i} \doteq \langle f, f \rangle_{R_i}$. It is easy to verify that R_i is a linear subspace of Ω_1 , and the topology induced by $\|\cdot\|_{R_i}$ is finer than that induced by $\|\cdot\|_{\Omega_1}$.

The RKHS of the multivariate process $(Z_t^{\{1\}}, \dots, Z_t^{\{k\}})$ is, by the independence of the $Z^{\{i\}}$'s, $R \doteq R_1 \times \dots \times R_k$ with the inner product

$$\langle (f_1, \dots, f_k), (g_1, \dots, g_k) \rangle_R \doteq \sum_{i=1}^k \langle f_i, g_i \rangle_{R_i}.$$

The *reproducing kernel property*, which is a straightforward consequence of the definition of the inner products, tells that within R , the functions can be evaluated by taking an inner product with a corresponding vector of covariance functions:

$$\langle (f_1, \dots, f_k), (\Gamma_1(t_1, \cdot), \dots, \Gamma_k(t_k, \cdot)) \rangle_R = \sum_{i=1}^k f_i(t_i). \tag{7}$$

The above construction can be further extended to the case that the component processes are dependent, as long as all joint distributions are Gaussian.

A large deviation principle for Gaussian measures in Banach space is given by the generalized Schilder's theorem (Bahadur and Zabell [6], see also [5,9]).

Theorem 1. *The function $I : \Omega \rightarrow \mathfrak{R} \cup \{\infty\}$,*

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

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

$$\text{for } F \text{ closed in } \Omega : \quad \limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P} \left(\frac{Z}{\sqrt{n}} \in F \right) \leq - \inf_{\omega \in F} I(\omega);$$

$$\text{for } G \text{ open in } \Omega : \quad \liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P} \left(\frac{Z}{\sqrt{n}} \in G \right) \geq - \inf_{\omega \in G} I(\omega).$$

Thus, the essential problem is to find a path ω that minimizes $I(\omega)$ in a given set B , or, equivalently, the norm $\|f\|_R$ in the set $B \cap R$. We call it the *most probable path* in that set. Intuitively, one can think of $e^{-I(\omega)}$ as something like the probability density of our infinite dimensional Gaussian measure, so that minimizing $I(\omega)$ corresponds to maximizing likelihood. In most cases, the most probable path is unique, but the examples in Section 4 show that non-unique paths may appear and even have interesting meaning as “phase transitions” of the queueing system.

The approach presented here was originally motivated by the generalized Schilder’s theorem [20]. However, our main interest is not in large deviations limits but in estimates that are applicable for whole distributions. It was shown in [2] by examples of ordinary queues that estimates of the type $\mathbb{P}(A) \approx \exp(-\inf_{\omega \in A} I(\omega))$ give indeed often a reasonable approximation of the whole queue length distribution, not only for tail behavior. On the other hand, note that it is problematic to even formulate large deviations limit theorems with Gaussian traffic, because the Gaussian character is already the result of another kind of limit procedure, the Central Limit Theorem.

3.2 Half-Space Approximations

Consider first the case of a simple queue. What can be said about the marginal distribution of Q_t ? Writing (cf. [17])

$$\{Q_t > x\} = \left\{ \sup_{s \leq t} \frac{Z_t - Z_s}{x + (c - m)(t - s)} > 1 \right\}$$

we see that this event is in fact of the form $\left\{ \sup_s Y_s^{(x,t)} > 1 \right\}$ for the *centered* Gaussian process $Y_s^{(x,t)} = (Z_t - Z_s)/(x + (c - m)(t - s))$. Thus, we encounter the very classical problem of estimating the distribution of the maximum of a centered Gaussian process. Consider the obvious lower bound

$$\mathbb{P}(Q_t > x) \geq \sup_{s \leq t} \mathbb{P} \left(Y_s^{(x,t)} > 1 \right) = \bar{\Phi} \left(\frac{x + (c - m)u^*}{\sqrt{v(u^*)}} \right) = \ell(x), \quad (8)$$

where $\bar{\Phi}$ is the residual distribution function of the standard normal distribution and $u^* > 0$ minimizes $(x + (c - m)u)^2/v(u)$ w.r.t. u . The value u^* has the important practical meaning of characterizing the *relevant timescale* of queues of length x .

Note the geometry of the set $\{Q_t > x\}$: it is the union over s of the sets $\{A(t - s) - c(t - s) > x\}$ which are half-spaces, and thus the complement of a convex set containing the origin. Let f^* be a most probable path in $\{Q_t > x\}$. The following proposition, which we formulate directly in the multiclass case, gives an explicit expression of f^* .




Fig. 1. The half-space $\{-Z_{-t^*} \geq x + (c - m)t^*\}$ is contained in the set $\{Q_0 \geq x\}$. For both sets, the closest point to origin is f^* .

Proposition 1. *Most probable path vectors f^* in the set $\{Q_0^{\{1, \dots, k\}} \geq x\}$ have the form*

$$-\frac{x + (c - m)t^*}{\sum_{i=1}^k v_i(t^*)} (\Gamma_1(-t^*, \cdot), \dots, \Gamma_k(-t^*, \cdot)),$$

where $t^* > 0$ minimizes the expression

$$h(t) = \frac{(x + (c - m)t)^2}{\sum_{i=1}^k v_i(t)}. \tag{9}$$

Proof. Note that

$$\begin{aligned} \{Q_0^{\{1, \dots, k\}} \geq x\} &= \bigcup_{s \leq 0} \{A^{\{1, \dots, k\}}(s, 0) - c(0 - s) \geq x\} \\ &= \bigcup_{s \leq 0} \{Z^{\{1, \dots, k\}}(s, 0) \geq x + (c - m)(0 - s)\}, \end{aligned}$$

and, by the reproducing kernel property,

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

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 $((x + (c - m)t) / \sum v_i(t))(\Gamma_1(-t, \cdot), \dots, \Gamma_k(-t, \cdot))$ with respect to $t > 0$.

Let $f^* \in R$ be a most probable path in a closed set $B \subset \Omega$ such that $f^* \neq 0$. We call the set

$$B^* \doteq \text{cl}_\Omega \{g \in R : \langle g - f^*, f^* \rangle_R \geq 0\},$$

where $\text{cl}_\Omega G$ denotes the closure of G in the topology of Ω , the *half-space approximation* of B . In particular, it is easy to see that

$$\{Q_0 \geq x\}^* = \{-Z_{-t^*} \geq x + (c - m)t^*\},$$

and the lower bound (8) is a consequence of the fact that in this case the half-space approximation is contained in the original set. See Figure 1.

It is worth of noting also that the most probable path vector in a set $\{A_t^{\{1, \dots, k\}} \geq y\}$, where $y > mt$, is in fact the conditional expectation

$$\mathbb{E} \left[(Z_s^{\{1\}}, \dots, Z_s^{\{k\}}) \mid A_t^{\{1, \dots, k\}} = y \right].$$

This is a consequence of the fact that the conditional distribution of a Gaussian vector w.r.t. a linear condition is Gaussian, and its expectation equals the point where the density is highest.

More accurate estimates take, in some way or other, the geometry of the set $\{Q_0 \geq x\}$ around f^* into account. For different methods, see the books by Adler [4] and Piterbarg [25]. An original geometric reasoning, after transforming the problem into Fourier space, was given in [18].

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 identifying these paths has an essential role in choosing a good approximation in the case of GPS and priority queues.

3.3 General Heuristic Approximations

Within logarithmic accuracy, the lower bound can be replaced by the still simpler approximate expression

$$\mathbb{P}(Q_t > x) \approx \exp \left(-\frac{(x + (c - m)u^*)^2}{2v(u^*)} \right), \quad (10)$$

which was called the *basic approximation* in [3]. Simulations of many cases indicate that the basic approximation may in fact be a general *upper bound* of $\mathbb{P}(Q_t > x)$, but no proof of this is known.

In all empirical studies and simulations one works in discrete time. The discrete time queue is always a little smaller than the corresponding continuous time queue. Indeed, if A is our continuous time model, the cumulative input process in discrete time is simply $(A_n)_{n \in \mathbb{Z}}$,

$$Q_n^{discr} = \sup_{m \leq n, m \in \mathbb{Z}} (A(m, n) - c(n - m)) \leq \sup_{s \leq n, s \in \mathbb{R}} (A(s, n) - c(n - s)) = Q_n^{cont}.$$

It was observed in [3] that one often gets fairly good approximations for a discrete time Gaussian queue Q^{discr} by multiplying the basic approximation by an appropriate constant p such that

$$p \lim_{x \rightarrow 0^+} \exp\left(-\frac{(x + (c - m)u_x^*)^2}{2v(u_x^*)}\right) \approx \mathbb{P}(Q_t^{discr} > 0).$$

A good heuristic approximation for the non-emptiness probability of a discrete time queue with time resolution δ is (see [3])

$$\mathbb{P}(Q_t^{discr} > 0) \approx 2\mathbb{P}(A_\delta > c\delta).$$

3.4 Approximations for GPS and Priority Queues

The structure of our method for getting estimates of queue length distributions in GPS and priority systems is the following. In order to get an approximation for $\{Q_0^{\{i\}} > x\}$, do

Step 1. Find the most probable path vector f^* of the event $\{Q_0^{\{1, \dots, k\}} > x\}$.

The path vector can be immediately written and plotted using Proposition 1.

Step 2. Check whether $Q_0^{\{1, \dots, k\} \setminus \{i\}}(f^*) = 0$. If yes, go to Step 3, otherwise go to Step 4.

Step 3. (Empty Buffer Approximation) f^* is the most probable path vector in $\{Q_0^{\{i\}} > x\}$; use the corresponding half-space approximation. Stop.

Step 4. (Rough Full Link Approximation) Find a certain f^{RFLA} , where the only positive queue is $Q_0^{\{i\}}$ (or the others are much smaller); use the half-space approximation corresponding to f^{RFLA} .

The Empty Buffer Approximation uses the true most probable path vector, and it can be considered as reliable as the simple queue estimates of Section 3.2. In the Rough Full Link Approximation, the path vector f^{RFLA} also is just a heuristic approximation of the true most probable path vector. Both approximations are discussed in more detail below.

The Empty Buffer Approximation. The idea of the Empty Buffer Approximation (EBA), first studied by Berger and Whitt [7,8], is that in a two-class priority queue, the total queue usually consists almost exclusively of lower class traffic, and therefore its distribution is a good approximation to that of the pure lower class queue. Our approach gives a straightforward method to check the applicability of EBA in any particular combination of Gaussian traffic streams. Our examples indicate that EBA is a very good principle in most practically interesting priority scenarios with Gaussian traffic.

The EBA is also often useful in the study of a GPS system. However, it is never sufficient, because the classes are in a symmetric position in GPS, and the distribution of at most one class can be estimated with EBA.

Whereas it may require some work to check analytically whether the most probable path vector producing joint queue x satisfies the EBA condition, an approximately similar condition is much easier: in the two-class priority case, just check the “rough EBA condition”

$$-A_{-t^*}^{\{1\}}(f^*) \leq ct^*. \tag{11}$$

This also leads to some interesting insight. Consider the priority system with two classes and assume, without restricting generality, that $m_1 = 0$. The condition 11 can be written as

$$\frac{x}{t^*} - m_2 \leq \frac{v_2(t^*)}{v_1(t^*)}c. \tag{12}$$

In particular, we see that (12) holds if $m_2 \geq x/t^*$ (note, however, that t^* depends on the other quantities). In the special case that v_1 is a multiple of v_2 , say $v_2(t) = av(t)$, $v_1(t) = bv(t)$, the condition becomes still simpler. Then t^* is independent of a and b , and we obtain the rather surprising result that when m_2 exceeds a certain threshold, then we are roughly in the EBA *irrespective* of the variance coefficients a and b ! For example, if both $Z^{(i)}$'s are fractional Brownian motions with same self-similarity parameter H , then $t^* = Hx/((1-H)(c-m_2))$, which gives the condition $m_2 \geq (1-H)c$. The higher H , the lower is the threshold for m_2 above which a typical large class 2 queue consists of class 2 traffic alone.

The Rough Full Link Approximation. Consider the case of two traffic classes. For priority queues this does not restrict generality (since we neglect the effect of negative traffic). For GPS queues, the idea below 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.

Consider a GPS system with weights μ_1 and μ_2 . In the two class case, the priority system is obtained as the special case $\mu_2 = 0$. Assume that we are interested in the number $\mathbb{P}\left(Q_0^{\{2\}} \geq x\right)$. As before, we first identify the most probable path pair f^* of $\left\{Q_0^{\{1,2\}} \geq x\right\}$. If $Q_0^{\{1\}}(f^*) = 0$, we can use the EBA, as discussed in Section 3.4. So assume that $Q_0^{\{1\}}(f^*) > 0$.

The idea of our approximation in the non-EBA case is that any superfluous queue buildup decreases the likelihood of our path pair. Since we are only requiring that $Q_0^{\{2\}}(\omega)$ be big, $Q_0^{\{1\}}(\omega)$ must be close to zero with the optimal ω .

Thus, a class 2 queue of size x is most easily made so that the role of class 1 is essentially to fill its quota (in the priority case, to fill the whole link) without making a queue, while class 2 fills its quota and additionally builds a queue of size x .

To make this condition still simpler, we reduce this behavior to the one-dimensional conditions

$$\begin{aligned} A^{\{1\}}(-t, 0) &= \mu_1 ct, \\ A^{\{2\}}(-t, 0) &= \mu_2 ct + x, \end{aligned} \tag{13}$$

write down the most probable path pair fulfilling this, and finally minimize their norm with respect to t . We call this procedure the Rough Full Link Approximation (RFLA).

It is again an easy Hilbert space exercise, similar to Proposition 1, to determine the most probable paths in RFLA (see [15]):

Proposition 2. *The most probable path pair f^{RFLA} satisfying (13) is of the form*

$$\begin{aligned} f^{\text{RFLA}}(\cdot) &= (f_1^{\text{RFLA}}(\cdot), f_2^{\text{RFLA}}(\cdot)) \\ &= \left(\frac{(\mu_1 c - m_1)t^*}{v_1(t^*)} \Gamma_1(t^*, \cdot), \frac{-x + (\mu_2 c - m_2)t^*}{v_2(t^*)} \Gamma_2(t^*, \cdot) \right), \end{aligned}$$

where $t^* < 0$ minimizes, w.r.t. t , the expression

$$\frac{(\mu_1 c - m_1)^2 t^2}{v_1(t)} + \frac{(x - (\mu_2 c - m_2)t)^2}{v_2(t)}. \tag{14}$$

In the case that both classes are Brownian motions (counterpart of Poisson processes), the RFLA gives the true most probable path pair in the non-EBA case. In general, however, the class 1 path in RFLA does not fill its quota over the whole interval $(-t^*, 0)$, thus part of class 2 traffic is “wasted”, and there is a small class 1 queue at time 0, whereas the class 2 queue remains correspondingly smaller than x .

Using the reproducing kernel property and the fact that evaluation at a time point is a continuous linear functional both in R and Ω , we see that the half-space corresponding to f^{RFLA} can be written as $E = \{Y \geq \|f^{\text{RFLA}}\|_R^2\}$, where

$$Y = \frac{(\mu_1 c - m_1)t^*}{v_1(t^*)} Z_{t^*}^{\{1\}} + \frac{x - (\mu_2 c - m_2)t^*}{v_2(t^*)} Z_{t^*}^{\{2\}}.$$

Thus, our RFLA approximation, which the simulations indeed indicate to be a lower bound, is

$$\begin{aligned} \mathbb{P}(Q_0^{\{2\}} \geq x) &\approx \mathbb{P}(E) \\ &= \bar{\Phi} \left(\sqrt{\frac{(\mu_1 c - m_1)^2 t^{*2}}{v_1(t^*)} + \frac{(x - (\mu_2 c - m_2)t^*)^2}{v_2(t^*)}} \right). \end{aligned} \tag{15}$$

In order to check accuracy of our estimates, we have compared them to the empirical measures calculated from simulations. The simulation traces were generated using an extension of random midpoint displacement algorithm (RMD_{mn}, see [21]). Many examples are included in the papers [3,15], and they show reasonable accuracy of the method. In particular, the “basic approximations” turn always out to be upper bounds and the probabilities of the half-space approximations lower bounds. In the present overview paper, we restrict to the following example taken from [15].

Example: two fBm traffic classes with same self-similarity parameter. Let us consider a GPS system with two classes with $v_i(t) = \sigma_i^2 t^{2H}$ for $i = 1, 2$, and the parameter H is any number in $(0, 1)$. In this case we can compute the above quantities analytically.

First, fix $x > 0$ and consider the total queue. We have (see, e.g., [3])

$$t^* = \frac{Hx}{(1-H)(c-m)}.$$

Second, the rough EBA criterion (cf. (12)) for estimating class 1 reads

$$(\mu_2 c - m_2)t^* \geq f_2^*(t^*) = (x + (c - m)t^*) \frac{\sigma_2^2}{\sigma_1^2 + \sigma_2^2}.$$

Substituting t^* , we obtain the criterion

$$\frac{(\mu_2 c - m_2)H}{c - m} \geq \frac{\sigma_2^2}{\sigma_1^2 + \sigma_2^2}. \quad (16)$$

Note that only the mean and service rates appear on left and only the variance coefficients on right. If (16) is satisfied, the “basic approximation” reads

$$\mathbb{P}\left(Q_0^{\{1\}} \geq x\right) \approx \mathbb{P}\left(Q_0^{\{1,2\}} \geq x\right) \approx \exp\left(-\frac{(c-m)^{2H}}{\sigma_1^2 + \sigma_2^2} \cdot \frac{x^{2-2H}}{2\kappa(H)^2}\right),$$

where $\kappa(H) = H^H(1-H)^{1-H}$.

Third, if (16) does not hold, we use the RFLA. The squared R -norm of the most probable path in the set

$$\left\{-A_{-t}^{\{1\}} \geq \mu_1 ct + x, -A_{-t}^{\{2\}} \geq \mu_2 ct\right\}$$

is

$$\frac{((\mu_1 c - m_1)t + x)^2}{\sigma_1^2 t^{2H}} + \frac{(\mu_2 c - m_2)^2}{\sigma_2^2} t^{2-2H}.$$

The minimum is obtained at $t^* = \eta x$, where η is the positive root of a quadratic equation:

$$\eta = \frac{b + \sqrt{b^2 + 4aH}}{2a}, \text{ where}$$

$$a = \left(\frac{(\mu_1 c - m_1)^2}{\sigma_1^2} + \frac{(\mu_2 c - m_2)^2}{\sigma_2^2}\right)(1-H), \quad b = \frac{(\mu_1 c - m_1)(2H-1)}{\sigma_1^2}.$$

The basic approximation of $\mathbb{P}(Q_0^{\{1\}} \geq x)$ is then

$$\begin{aligned} & \mathbb{P}(Q_0^{\{1\}} \geq x) \\ & \approx \exp\left(-\frac{1}{2}\left(\frac{((\mu_1 c - m_1)\eta + 1)^2}{\sigma_1^2 \eta^{2H}} + \frac{(\mu_2 c - m_2)^2}{\sigma_2^2} \eta^{2-2H}\right) x^{2-2H}\right). \end{aligned} \tag{17}$$

Simulations indicate that the approximations of this section work quite well — see [14,13,15].

4 “Phase Transitions” of Typical Queues

Even for simple queues, the most probable paths need not be unique. Nice examples of this were found by P. Mannersalo by superposing a periodic source and a fBm source [3,15]. A kind of phase transition was observed: typical small queues were caused by the periodical fluctuation of the periodic traffic, whereas typical long queues were caused by sustained heightened activity of the fBm traffic. (Cf. also [12].)

Another and, most importantly, non-artificial example was encountered by Pazhyannur and Fleming [24]. They studied a queue with input consisting of periodic coded voice traffic, modelled as follows. A source transmits with period d and uniformly distributed phase U . Volume in i th period is X_i . The X_i ’s can be strongly dependent. There are n i.i.d. sources. See Figure 2.




Fig. 2. The structure of vocoder traffic in [24]. Each source transmits periodically bursts whose sizes are random but correlated. The X_i ’s in the picture come from the same source.

Assuming that the number of sources is large enough for Gaussian modelling, our technique can be applied in a straightforward way. We only need to compute $v(t)$ for a single source — using a mathematical computer tool, the rest follows “according to the recipe”. Denote the phase of our source by U and choose, for simplicity, $d = 1$. Then

$$\begin{aligned} A_t &= \sum_{i=1}^{\lfloor t \rfloor} X_i + 1_{\{U < t - \lfloor t \rfloor\}} X_{\lfloor t \rfloor + 1} \\ v(t) &= t \text{Var}(X_0) \end{aligned}$$

$$\begin{aligned}
 &+2 \sum_{k=1}^{\lfloor t \rfloor} (t-k) \text{Cov}(X_0, X_k) + (t - \lfloor t \rfloor)(1 - (t - \lfloor t \rfloor))(\mathbb{E}\{X_0\})^2 \\
 \Gamma(s, t) &= \frac{1}{2}(v(|s|) + v(|t|) - v(|s - t|))
 \end{aligned}$$

Consider, as an example, the case $m = 0$, $\text{Cov}(X_0, X_k) = \rho^k$, $\rho = 0.9$. Figure 3 shows a clear bend in the complementary distribution function of the queue length (resembling the shift from the “cell scale queue” to the “burst scale queue” in many ATM analyses — see, e.g., [22]). What happens when the queue size increases from 0.3 to 0.4?




Fig. 3. Estimate of $\log_{10} \mathbb{P}(V > x)$.

Look first at the function $h(t) = (x + t)^2/v(t)$, which has to be minimized with respect to t . For $x = 0.3$ or smaller, we have $t^* \approx x$, whereas for $x = 0.4$, $t^* \approx 4$. Somewhere between 0.3 and 0.4 is a value of $x = x_0$ where the two local minima are equal. As a function of x , t^* makes big jump at x_0 .




Fig. 4. Plot of the function $\frac{(x + t)^2}{v(t)}$. Left: $x = 0.3$. Right: $x = 0.4$.

Finally, the most probable paths shows that there is a very clear difference between typical queues of sizes 0.3 and 0.4. In the former, the queue is caused

only by the bursts from different users, which are independent. In the latter, the busy period is larger than the period of the sources, which has the effect that the strong correlations between bursts of each source have become dominant, and the distribution tail decreases much slower than it did for small x 's. Pazhyannur and Fleming discovered this queue behavior originally using more traditional heavy traffic approximations, but our method added an immediate visual insight which agreed with their interpretation. Moreover, they found that the Gaussian approximations were also quantitatively quite good.




Fig. 5. Most probable queue path. Left: $x = 0.3$. Right: $x = 0.4$.

5 A Simple Model for Bandwidth Allocation by Prediction

Our last example analyses the performance of a queue whose service capacity is dynamically adjusted according to predicted demand, with a fixed prediction delay. The following setup is probably the simplest possible model for that kind of system.

Let A_t again be a Gaussian traffic process with parameters m and $v(t)$. Assume that instead of a fixed service rate, the service capacity is allocated dynamically with a delay Δ , with a relative surplus capacity ϵ . That is, we define the cumulative service process as

$$C_t \doteq (1 + \epsilon)(A_{t-\Delta} - A_{-\Delta}) \tag{18}$$

(the last term is included in order to have $C_0 = 0$). The queue length process is

$$Q_t = \sup_{s \leq t} (A(s, t) - C(s, t)) \\ \stackrel{\mathcal{D}}{=} \sup_{t \geq 0} (U_t - \epsilon mt),$$

where $U_t = Z_t - (1 + \epsilon)(Z_{t+\Delta} - Z_{\Delta})$. A straightforward computation gives

$$\text{Var}(U_t) = (1 + (1 + \epsilon)^2)v(t) - (1 + \epsilon)(v(t - \Delta) + v(t + \Delta)) + 2(1 + \epsilon)v(\Delta).$$

In the space R we have

$$f(t) - (1 + \epsilon)(f(t + \Delta) - f(\Delta)) = \langle f, \Gamma(t, \cdot) - (1 + \epsilon)(\Gamma(t + \Delta, \cdot) - \Gamma(\Delta, \cdot)) \rangle_R.$$

Thus, by our general method, the most probable path of Z creating a queue of size x at time 0 is

$$f_x^*(s) = -\frac{x + \epsilon mt^*}{\text{Var}(U_{t^*})}(\Gamma(-t^*, s) - (1 + \epsilon)(\Gamma(-t^* \Delta, s) - \Gamma(-\Delta, s))),$$

where $t = t^* > 0$ minimizes

$$\frac{(x + \epsilon mt)^2}{\text{Var}(U_t)}.$$

In fact, the delay in such a system is bounded by Δ . The delay of a “fluid molecule” entering the system at time t can be expressed as

$$D_t \doteq \inf \{ \tau : C(t, t + \tau) \geq Q_t \}.$$

Now,

$$\begin{aligned} Q_t - C(t, t + \tau) &= \sup_{s \leq t} (A(s, t) - (1 + \epsilon)A(s - \Delta, t - \Delta)) \\ &\quad - (1 + \epsilon)A(t - \Delta, t - \Delta + \tau) \\ &= \sup_{s \leq t} (A(s, t) - (1 + \epsilon)A(s - \Delta, t - \Delta + \tau)) \leq 0 \end{aligned}$$

for $\tau \geq \Delta$, assuming that A_t is nondecreasing (which does not hold strictly for a Gaussian traffic model). (I thank P. Mannersalo for this insight.)

As an example, let us look at some paths in the case of fBm input $A_t = mt + \sigma Z_t$, where Z is a normalized fBm with self-similarity parameter H . The figures below were made with $\epsilon = 0.1$, $\Delta = 1$, $m = 3$, $\sigma^2 = 1$, and $H = 0.75$.

Figure 6 compares the dynamically varied service with fixed service rate and same 10% overallocation. It is no surprise that very big queues arise when such a high load is offered to a fixed capacity server, whereas the queue remains essentially bounded in the former case (remember that the delays are strictly bounded). Figure 7 shows lower bound estimates of the complementary distribution functions. Indeed, the distribution tail of the dynamically served queue decreases very fast (faster than exponentially).

Figure 7 shows the most probable paths of the input rate and the queue of size 4. Note how cleverly our system makes its big (by its scale) queues: in order to fool the prediction, the input is first very slow and then, when the control cannot react any more, it suddenly speeds up. The queue path also has a noteworthy feature: after an input peak, the typical queue first decreases quickly, but then shifts to much slower decrease, whose slope corresponds to the overhead ϵ .




Fig. 6. Queue length processes of a system with prediction based dynamic allocation with $\epsilon = 0.1$, $\Delta = 1$ (left), and a system with fixed service capacity $(1+\epsilon)m$ (right). The input processes are identical discrete time fBm traces with $m = 3$, $\sigma^2 = 1$, $H = 0.75$.




Fig. 7. Queue length distribution lower bounds $\log_{10} \ell(x)$ (see (8)) for a system with prediction based dynamic allocation with $\epsilon = 0.1$, $\Delta = 1$ (squares), and a system with fixed service capacity $(1+\epsilon)m$ (stars). The input processes are fBm with $m = 3$, $\sigma^2 = 1$, $H = 0.75$.




Fig. 8. The most probable path with queue size $x = 4$ in a system with prediction based dynamic allocation with $\epsilon = 0.1$, $\Delta = 1$. The input process is fBm with $m = 3$, $\sigma^2 = 1$, $H = 0.75$. Left: rate. Right: queue length.

6 Conclusion

We have presented a straightforward method for studying various queueing systems with general Gaussian input traffic. These included priority queues, two-class GPS queues, and an example of dynamic server capacity allocation.

Using any advanced mathematical tool, it is possible to build expert systems, which make the analyses in this paper half or fully automatic once the parameters are given. In particular, the traffic in each class is described simply with mean rate and the cumulative variance function.

The novel theoretical aspect in this work is that we are looking for approximations and bounds in a Gaussian space — not large deviation theorems, which at least “officially” tell only about certain logarithmic limits. Although most of our quantitative estimates are more or less heuristic, we hope that this new point of view to queueing phenomena will prove fruitful in rigorous mathematics also. One of the key challenges may then be understanding the geometry of the threshold exceedance set in the neighborhood of the most probable path.

References

1. R.G. Addie. On weak convergence of long range dependent traffic processes. *Journal of Statistical Planning and Inference*, 80(1-2):155–171, 1999.
2. R.G. Addie, P. Mannersalo, and I. Norros. Performance formulae for queues with Gaussian input. In P. Key and D. Smith, editors, *Teletraffic Engineering in a Competitive World. Proceedings of the International Teletraffic Congress — ITC-16*, pages 1169–1178, Edinburgh, UK, 1999. Elsevier.
3. R.G. Addie, P. Mannersalo, and I. Norros. Most probable paths and performance formulae for buffers with Gaussian input traffic. To appear in *European Transactions on Telecommunications*, 2002.
4. R.J. Adler. *An Introduction to Continuity, Extrema, and Related Topics for General Gaussian Processes*, volume 12 of *Lecture Notes-Monograph Series*. Institute of Mathematical Statistics, 1990.
5. R. Azencott. *Ecole d’Eté de Probabilités de Saint-Flour VII-1978*, chapter Grandes deviations et applications, pages 1–176. Number 774 in *Lecture notes in Mathematics*. Springer, Berlin, 1980.
6. R.R. Bahadur and S.L. Zabell. Large deviations of the sample mean in general vector spaces. *Ann. Prob.*, 7(4):587–621, 1979.
7. A.W. Berger and W. Whitt. Effective bandwidths with priorities. *IEEE/ACM Transactions on Networking*, 6(4), 1998.
8. A.W. Berger and W. Whitt. Extending the effective bandwidth concept to networks with priority classes. *IEEE Communications Magazine*, August 1998.
9. J.-D. Deuschel and D.W. Stroock. *Large Deviations*. Academic Press, Boston, 1989.
10. J. Kilpi and I. Norros. Testing the Gaussian character of access network traffic. Technical Report COST279TD(01)03, COST, 2001. Available from <http://www.vtt.fi/tte/projects/cost279/>.
11. W.E. Leland, M.S. Taqqu, W. Willinger, and D.V. Wilson. On the self-similar nature of Ethernet traffic (extended version). *IEEE/ACM Transactions on Networking*, 2(1):1–15, February 1994.

12. M.R.H. Mandjes and J.H. Kim. An analysis of the phase transition phenomenon in packet networks. To appear in *Adv. or J. of Applied Probability*.
13. P. Mannersalo and I. Norros. GPS schedulers and Gaussian traffic. *Infocom 2002*, New York.
14. P. Mannersalo and I. Norros. Gaussian priority queues. In *Proceedings of ITC 17*. Elsevier, 2001.
15. P. Mannersalo and I. Norros. A most probable path approach to queueing systems with general Gaussian input. *Computer Networks*, 2002. To appear.
16. L. Massoulié. Large deviations estimates for polling and weighted fair queueing service systems. *Adv. Perf. Anal.*, 2(2):103–127, 1999.
17. L. Massoulié and A. Simonian. Large buffer asymptotics for the queue with FBM input. *J. Appl. Prob.*, 36(3):894–906, 1999.
18. O. Narayan. Exact asymptotic queue length distribution for fractional Brownian traffic. *Advances in Performance Analysis*, 1:39–63, 1998.
19. I. Norros. A storage model with self-similar input. *Queueing Systems*, 16:387–396, 1994.
20. I. Norros. Busy periods of fractional Brownian storage: a large deviations approach. *Adv. Perf. Anal.*, 2:1–19, 1999.
21. I. Norros, P. Mannersalo, and J.L. Wang. Simulation of fractional Brownian motion with conditionalized random midpoint displacement. *Adv. Performance Anal.*, 2:77–101, 1999.
22. I. Norros, J.W. Roberts, A. Simonian, and J.T. Virtamo. The superposition of variable bit rate sources in an ATM multiplexer. *IEEE JSAC*, 9(3):378–387, April 1991.
23. 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*, 1(3):344–357, 1993.
24. R.S. Pazhyannur and P. Fleming. Asymptotic results for voice delay in packet networks. In *Vehicular Technology Conference / Fall*, 2001.
25. V.I. Piterbarg. *Asymptotic Methods in the Theory of Gaussian Processes and Fields*. American Mathematical Society, 1996.
26. P. Tran-Gia and N. Vicari, editors. *Impacts of new services on the architecture and performance of broadband networks. COST 257 Final Report*. compuTEAM Würzburg, 2000. <http://nero.informatik.uni-wuerzburg.de/cost/Final/>.