Available online at www.sciencedirect.com
ScienceDirect
Computer Networks 51 (2007) 781–797
Computer Networks
www.elsevier.com/locate/comnet

Gaussian tandem queues with an application to dimensioning of switch fabric interfaces

Michel Mandjes [a,b,1], Petteri Mannersalo [c,*,2], Ilkka Norros [c]

[a] CWI, P.O. Box 94079, NL-1090 GB Amsterdam, The Netherlands
[b] University of Amsterdam, Korteweg–de Vries Institute for Mathematics, Plantage Muidergracht 24, NL-1018 TV Amsterdam, The Netherlands
[c] VTT, P.O. Box 1000, FI-02044 VTT, Finland

Received 22 June 2005; received in revised form 4 June 2006; accepted 20 June 2006
Available online 17 July 2006

Responsible Editor: E. Altman

Abstract

Tandem systems are seen in many places and at various hierarchical levels in high-speed communication networks, from router architectures to protocol stacks. If the traffic fed into the system is generated by independent or weakly dependent sources and the smallest relevant time scale is not too fine, the central limit theorem suggests that the input traffic is (close to) Gaussian.

This paper considers tandem queues fed by Gaussian processes with stationary increments. Relying on the generalized version of Schilder's sample-path large-deviations theorem, we derive the many-sources asymptotics of the overflow probabilities in the second queue; 'Schilder' reduces this problem to finding the most probable path along which the second queue reaches overflow. The general form of these paths is described by recently obtained results on infinite-intersections of events in Gaussian processes; for the special cases of fractional Brownian motion and integrated Ornstein–Uhlenbeck input, the most probable path can be explicitly determined, as well as the corresponding exponential decay rate.

As the computation of the decay rate is numerically involved, we introduce an explicit approximation ('rough full-link approximation'). Based on this approximation, we propose performance formulae for network provisioning purposes. Simulation is used to assess the accuracy of the formulae. As an example, we show how the methods can be applied to dimensioning the interface between a line card and a switch fabric.

Keywords: Tandem queue; Gaussian process; Large deviations; Performance formulae

1. Introduction

The tandem queue is one of the canonical 'building blocks' in high-speed communication networks: they are encountered in many places and at various hierarchical levels, ranging from router architectures

* Corresponding author.
E-mail addresses: Michel.Mandjes@cwi.nl (M. Mandjes), Petteri.Mannersalo@vtt.fi (P. Mannersalo), Ilkka.Norros@vtt.fi (I. Norros).
[1] M. Mandjes is also with EURANDOM, Eindhoven, The Netherlands.
[2] Work done while P. Mannersalo was on leave at CWI.

1389-1286/$ - see front matter
doi:10.1016/j.comnet.2006.06.006

to protocol stacks. This fact has motivated the development of queueing-theoretic methods for analyzing their performance (in terms of loss, delay, throughput, etc.). Although pure tandem queues (without cross-traffic) may seem to be fairly simple systems, they are quite difficult to handle analytically. The main reason for this lies in the fact that the statistical properties of the traffic streams change – usually in a fairly non-trivial way – when traversing network nodes, the so-called 'shaping effect'.

In this paper we develop an analytic approach to the large deviations of a tandem queue fed by Gaussian traffic with arbitrary correlation structure and demonstrate it both with bursty, long-range dependent traffic (fractional Brownian motion, fBm) and with a smooth and short-range dependent traffic model (integrated Ornstein–Uhlenbeck, iOU). In particular, we identify the most probable paths leading to a high level of the second queue.

As an example showing in what type of applications our techniques are potentially useful, we consider the problem of dimensioning the interface between a line card and a switch fabric, see e.g. [29]. This is an optimization problem that arises from the facts that (i) the switch fabric speeds are much higher than line speeds, (ii) ultra-fast buffer memory is much more expensive than moderately fast memory, and (iii) the interface speed is a freely designable parameter. Imposing some natural cost structure (where the cost is convex in the speed of the interface and, given the speed, linear in both buffer sizes), the dimensioning can be done by finding the cost-minimizing parameters. This procedure is illustrated by an appealing numerical example that suggests that with higher traffic load one should rather use a slower interface speed and larger buffer than the contrary.

Another, more generic, question that could be analyzed relying on the framework of this paper, relates to the shaping effect mentioned above. A traffic stream can be made more 'benign' by sending it through a network node, in the sense that this makes the stream 'smoother', and hence easier to handle at downstream nodes. A trivial observation in this respect is that the peak rate of the output stream is upper bounded by the link speed of the queue, but a more precise quantification of the shaping effect (i.e., the decrease in 'burstiness' of the stream due to traversing the node) is usually hard. By using the analysis presented in this paper, however, we could get an analytic handle on this issue.

Gaussian traffic and Gaussian tandem queues. Assuming stationarity, the Gaussian traffic model says that the amount of traffic arriving in an arbitrary interval of any length t has a Gaussian distribution $N(\mu t, v(t))$. Our main reason for focusing on a Gaussian traffic model is that despite its abstract, highly idealized character it is the simplest model that allows arbitrary correlation structures, long-range dependence in particular. (Long-range dependence is most succinctly expressed in terms of the variance of the traffic arriving in an interval of length t, which is proportional to t^{2H} over a wide range of values of t. The parameter H is referred to as the Hurst parameter [14] and typically takes values in the range 0.7–0.9.) On the other hand, the choice of a Gaussian model is far from arbitrary, since the central limit theorem suggests that traffic on communication links will become closer to Gaussian as more independent sources add their contribution [1]; see also [13].

We define a Gaussian tandem queue as follows (for ease we restrict ourselves to a two-node system, but the argument extends to tandems of any size). Denote by $A(s,t)$ the traffic arriving in $[s,t)$. To avoid trivialities, assume that the (constant) link speed of queue 1 (c_1) is larger than the link speed of queue 2 (c_2). Queue 1 is just a normal queue with Gaussian input (see [3]), and its size at time 0 is given by the standard formula $\sup_{t>0}(A(-t,0) - c_1 t)$. For the second queue, a 'reduction principle' applies: the *total* queue length at time 0 is given by $\sup_{t>0}(A(-t,0) - c_2 t)$ (see, e.g., [4,11]), so that we can define

$$Q_2 = \sup_{t>0}(A(-t,0) - c_2 t) - \sup_{t>0}(A(-t,0) - c_1 t). \tag{1}$$

Note that although Gaussian models have the anomaly that they allow negative input traffic, the fact that $c_1 > c_2$ implies that Q_2 is always non-negative.

Large deviations in the many-sources regime. Since exact analysis of Gaussian queues is possible only in a few special cases, we resort to asymptotic regimes. In this paper, we assume that n i.i.d. Gaussian sources feed into the queueing system, where the service rates of the queues as well as the buffer thresholds are scaled by n, too. We now let n go to infinity; the resulting framework is often referred to as the *many-sources* scaling, see e.g. [30]. A vast body of results exists for single FIFO queues under this scaling. Most notably, under very mild conditions on the source behavior, it is possible to calculate the *exponential* decay rate of the probability $p_n(b,c)$ that