On the Build-Up of Large Queues in a Queueing Model with Fractional Brownian Motion Input

Neil O'Connell, Gregorio Procissi*
Basic Research Institute in the Mathematical Sciences
HP Laboratories Bristol
HPL-BRIMS-98-18
August, 1998

E-mail: noc@hplb.hpl.hp.com
procissi@edison.iet.unipi.it

large deviations,
moderate deviations,
Schïlder's theorem,
calculus of
variations

We analyse the way in which large queues build up in the single-server fractional Brownian motion queueing model. The large deviations problem for the queue-length process can be rephrased as a moderate deviations problem for the underlying white noise. This framework allows us to obtain not only an asymptotic expression for the probability of overflow, but also the most likely path followed by the queue-length process to reach the overflow level and prediction of post-overflow behaviour. The model we consider has stationary increments: there is also a non-stationary version of fractional Brownian motion, introduced by Lévy, which formed the basis for a similar study by Chang, Yao and Zajic [9]. We compare our results with theirs, and illustrate the essential differences between the two models.

*Department of Information Engineering, University of Pisa, Pisa, Italy

On the build-up of large queues in a queueing model with fractional Brownian motion input

Neil O'Connell*and Gregorio Procissi†

BRIMS, Technical Report HPL-BRIMS-98-18

July 31, 1998

Abstract

We analyse the way in which large queues build up in the single-server fractional Brownian motion queueing model. The large deviations problem for the queue-length process can be rephrased as a moderate deviations problem for the underlying white noise. This framework allows us to obtain not only an asymptotic expression for the probability of overflow, but also the most likely path followed by the queue-length process to reach the overflow level and prediction of post-overflow behaviour. The model we consider has stationary increments: there is a also non-stationary version of fractional Brownian motion, introduced by Lévy, which formed the basis for a similar study by Chang, Yao and Zajic [9]. We compare our results with theirs, and illustrate the essential differences between the two models.

Keywords

Large Deviations, Moderate Deviations, Schïlder's Theorem, Calculus of Variations.

*Neil O'Connell is with Hewlett-Packard's Basic Research Institute in the Mathematical Sciences (BRIMS) at Hewlett-Packard Labs, Bristol, UK. E-mail: noc@hplb.hpl.hp.com.

†Gregorio Procissi is with NetGroup, Department of Information Engineering, University of Pisa, Italy. Research carried out while on a student placement at BRIMS. E-mail: procissi@edison.iet.unipi.it.

1 Introduction

Recent studies [1] on broadband networks suggest that real traffic exhibits Long Range Dependence (LRD). This would imply a hyperbolic decay of autocorrelations that cannot be parsimoniously captured using traditional Markovian models. Various LRD traffic models have been proposed. The canonical model [8] is based on *fractional Brownian motion* (FBM). This model has been widely adopted for its parsimonious structure, as it depends on just three parameters: mean, variance and Hurst parameter. The Hurst parameter reflects the degree of LRD.

In this paper we focus our attention on this FBM process and, in particular, the variational problem associated with *how* the process reaches a given (high) level, conditional on the event that it *does*. We apply this to the associated queueing model to determine analytically *the most likely path to overflow and most likely bevahiour thereafter.* We also deduce known results on the asymptotics of overflow probabilities, previously obtained by less informative methods [6].

The FBM model we consider has stationary increments. Chang, Yao and Zajic [9] present a similar study to ours, which is based on a version of FBM which does not have stationary increments. The two models exhibit quite different behaviour and it is interesting to compare: this is the topic of Section VII.

2 Representation for scaled FBM

We begin with the following representation for FBM, due to Mandelbrot and Van Ness [2], which expresses the FBM process as a functional of an underlying *Brownian motion* (BM) $B(t)$ (see also [3]):

$$B^H(t) = \frac{\sigma}{C_1(H)} \int_{-\infty}^{t} f(t,y)dB(y) \quad t \geq 0, \tag{1}$$

where $0 < H < 1$, $\sigma > 0$ and

$$f(t,y) = (t-y)^\alpha - (-y)^\alpha 1_{(-\infty,0)}(y), \tag{2}$$

$$\alpha = H - \frac{1}{2}, \tag{3}$$

$$C_1(H) = \left\{ \int_0^\infty ((1+y)^\alpha - y^\alpha)^2 \, dy + \frac{1}{2H} \right\}^{1/2}. \tag{4}$$

Note that the underlying Brownian motion is indexed by the entire real line; this is constructed by setting $B(0) = 0$ and running two independent Brownian motions, one forwards, and one backwards in time. We consider the following scaled versions of BM and FBM:

$$B_t^\beta(s) = \frac{B(st)}{t^\beta} \qquad \text{with } 1/2 < \beta < 1, \tag{5}$$

and

$$B_t^H(s) = \frac{B^H(st)}{t}. \tag{6}$$

The starting point is to find the correct scaling β such that the Mandelbrot-Van Ness representation holds for the scaled processes, that is:

$$B_t^H(s) = \frac{\sigma}{C_1(H)} \int_\infty^s f(s,y) dB_t^\beta(y); \tag{7}$$

by simple calculations we obtain $\beta = 1 - \alpha = \frac{3}{2} - H$.

3 The basic variational problem

In this section we determine the most likely path of the scaled FBM conditioned to reach a given value x at time τ.

By (7) and a moderate deviations version of Schïlder's theorem [4] we have[1]

$$\begin{aligned} \lim_{t\to\infty} \frac{1}{t^\gamma} P(B_t^H(\tau) \approx x) &= \lim_{t\to\infty} \frac{1}{t^\gamma} P\left(B_t^\beta(s) \in A\right) \\ &= -\inf_{\phi \in A} I(\phi), \end{aligned} \tag{8}$$

with, as shown in [6], $\gamma = 1 - 2\alpha = 2 - 2H$, and where

$$I(\phi) = \frac{1}{2} \int_{-\infty}^\tau \dot{\phi}^2(s) ds \tag{9}$$

[1]We have chosen to argue at a formal level to avoid cumbersome technical details; we have no doubt that our conclusions are correct and can be made rigorous by starting with an appropriate topology for the Brownian motion moderate deviations principle, such as the topology used in the version of Shïlder's theorem presented in [5].

and

$$A = \left\{ \phi \quad : \quad \frac{\sigma}{C_1(H)} \int_{-\infty}^{\tau} f(\tau, s)\dot{\phi}(s)ds = x \right\}. \tag{10}$$

Now, setting:

$$\dot{g}(s) = f(\tau, s)\dot{\phi}(s) \tag{11}$$

the problem is to minimize the integral:

$$\frac{1}{2}\int_{-\infty}^{\tau} \dot{\phi}^2(s)ds = \frac{1}{2}\int_{-\infty}^{\tau} \left[\frac{\dot{g}(s)}{f(\tau, s)}\right]^2 ds \tag{12}$$

subject the conditions:

$$g(-\infty) = 0, \tag{13}$$

$$g(\tau) = \frac{xC_1(H)}{\sigma}. \tag{14}$$

This is a classical problem (see, for example, [7]); we solve the corresponding *Euler equation* to obtain:

$$\dot{\phi}(s) = \frac{x}{\sigma\tau^{2H}C_1(H)} f(\tau, s). \tag{15}$$

It follows that the *most likely path* of the scaled FBM has the expression:

$$\phi_t^H(s) = \frac{x}{C_1^2(H)\tau^{2H}} \int_{-\infty}^{s} f(s, y)f(\tau, y)dy \quad 0 \le s \le \tau. \tag{16}$$

Furthermore, by (8), we can evaluate the cost for the most likely path of the underlying Brownian motion to obtain:

$$-\inf_{\phi \in A} I(\phi) = -\frac{1}{2\sigma^2}\frac{x^2}{\tau^{2H}}. \tag{17}$$

4 Predicting Future Behaviour

Once the process $B_t^H(s)$ has reached the level x at time τ, we may wonder what is the most likely trajectory it will follow thereafter. One can again argue formally that, at this scaling, the future behaviour can be determined by setting $\dot{\phi}(s) = 0$ for $s > \tau$ (this is mean behaviour for the underlying Brownian motion and the law of large numbers is in effect).

Thus we can write the following expression for the most likely path of scaled FBM, conditional on the event that it reaches the level x at time τ:

$$\phi_t^H(s) = \begin{cases} \dfrac{x}{C_1^2(H)\tau^{2H}} \displaystyle\int_{-\infty}^{s} f(s,y) f(\tau,y) dy & s \leq \tau \\ \dfrac{x}{C_1^2(H)\tau^{2H}} \displaystyle\int_{-\infty}^{\tau} f(s,y) f(\tau,y) dy & s > \tau \end{cases} \tag{18}$$

5 Consequences for the Queuing Model

In this section we exploit the previous result in the analysis of the FIFO general single-server queue with infinite buffer size. In particular we obtain an asymptotic expression for the complementary probability $P(Q > b)$, typically used as an estimator for the overflow probability when the buffer has finite capacity b, as well as the most likely way and the most likely time at which the queue-length reachs a given critical level. As in [8], we suppose that the workload process (the difference between the total amount of work brought to the queue in the last t time units and the total amount of work that can be processed by the queue in the same time interval), is given by

$$W(t) = B^H(t) - \mu t \quad \text{for } t \geq 0, \tag{19}$$

where $B^H(t)$ is FBM and μ is a positive drift parameter. As is well known (see for example [6]) the current queue-length (which necessarily realises the equilibrium queue-length distribution) is given by:

$$Q = \sup_t W(t); \tag{20}$$

applying the *Principle of the Largest Term* (see for example [6]) with the substitution $t = \tau b$ we obtain:

$$\begin{aligned} \lim_{b\to\infty} \frac{1}{b^\gamma} \log P(Q > b) &= \lim_{b\to\infty} \frac{1}{b^\gamma} \log P\left(\sup_t \{B^H(t) - \mu t\} > b\right) \\ &= \lim_{b\to\infty} \frac{1}{b^\gamma} \log P\left(\sup_\tau \{B^H(\tau b) - \mu \tau b\} > b\right) \\ &\sim \sup_\tau \lim_{b\to\infty} \frac{1}{b^\gamma} \log P\left(B_b^H(\tau) > 1 + \mu\tau\right) \\ &= \sup_\tau - \inf_{\phi\in A} I(\phi), \end{aligned} \tag{21}$$

where, as before,

$$I(\phi) = \frac{1}{2} \int_{-\infty}^{\tau} \dot{\phi}^2(s) ds, \tag{22}$$

and

$$A = \left\{ \phi \ : \ \frac{\sigma}{C_1(H)} \int_{-\infty}^{\tau} f(\tau, s)\dot{\phi}(s)ds = 1 + \mu\tau \right\}. \tag{23}$$

Now, calling ϕ^* the most likely path of underlying scaled BM as in the previous section (and which depends on τ), we can write:

$$\begin{aligned} \lim_{b\to\infty} \frac{1}{b^\gamma} \log P(Q > b) &= -\inf_\tau I(\phi^*) \\ &= (\text{from (17) where } x = 1 + \mu\tau) \\ &= -\frac{1}{2\sigma^2} \left(\frac{\mu}{H}\right)^{2H} \frac{1}{(1-H)^{2-2H}}, \end{aligned} \tag{24}$$

agreeing with [8] and [6]. Furthermore, we find that:

$$\begin{aligned} \tau^* &= \arg\inf\{I(\phi^*)\} \\ &= \frac{H}{\mu(1-H)} \end{aligned} \tag{25}$$

is the *most likely time* at which the queue-length reachs the level b. The normalized *most likely queue-length trajectory* can be written as:

$$Q_b(s) = \max\left\{0, \phi_b^H(s) - \mu s\right\}, \tag{26}$$

where $\phi_b^H(s)$ has been defined in (18) with $t = b$, $x = 1 + \mu s$ and $\tau = \tau^*$.

Figures 1, 2 and 3, show with solid lines the behaviour of (26) for different values of H and μ.

6 How Does the Most Likely Path Depend on the Parameters?

The queuing model analysed depends on three parameters: the Hurst parameter H, μ and σ. We want to investigate how the shape of the most likely path depends on these three parameters. First of all note that $\phi_t^H(s)$ does not depend on σ (note however that $I(\phi^H)$, the cost, does). On the other hand $\phi_t^H(s)$ (and thus $Q_b(s)$) strictly depends on H and that dependency is evident in the pictures. Roughly speaking we observe that the closer H is to 0.5 (the BM case), the closer the aspect of the path is to a straight line; the most likely path becomes more and more 'non-linear' as H increases. Finally it is also interesting to point out that the shape of the path does not depend on μ and this can be readily verified by simple calculations using (18) and (25).

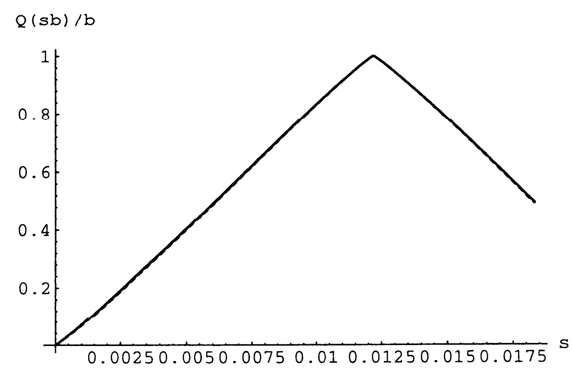

Figure 1: Normalized most likely trajectory followed by the queue length for $H = 0.55$ and $\mu = 100$ with FBM (solid line) and Levy's version of FBM (dashed line) input processes.

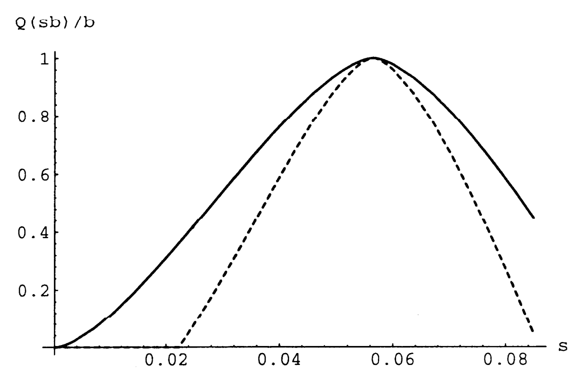

Figure 2: Normalized most likely trajectory followed by the queue length for $H = 0.85$ and $\mu = 100$ with FBM (solid line) and Levy's version of FBM (dashed line) input processes.

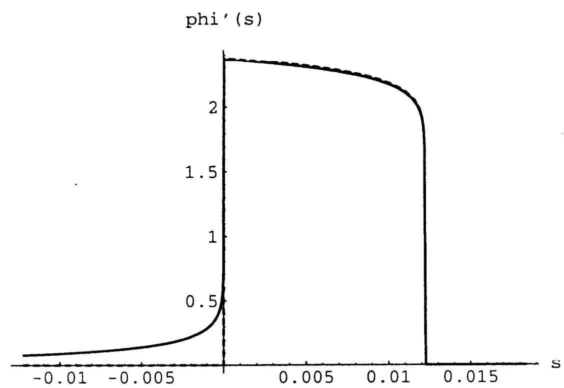

Figure 3: Derivative of underlying Brownian Motion for FBM (solid line) and Levy's version of FBM (dashed line) when $H = 0.55$ and $\mu = 100$.

7 Camparison with the Levy version of FBM

An interesting comparison can be made when, instead of the version of FBM with stationary increments, we consider the 'Lévy version', which formed the basis for a similar study by Chang, Yao and Zajic [9]. Using the same notation as in Section I, the Lévy version of FBM (see also [2]) can be represented as:

$$B^L(t) = \frac{\sigma}{C_1(H)} \int_0^t (t-y)^\alpha dB(y) \quad t \geq 0, \tag{27}$$

where now, for fair comparison, we set

$$C_1(H)^2 \cdot 2H = 1 \tag{28}$$

so that $Var B^L(t) = \sigma^2 t^{2H}$.

Following the same approach and solving an analogous variational problem as in Sections III and IV, we can write the expression for the most likely path of Levy version of FBM, conditional on the event that it reaches the

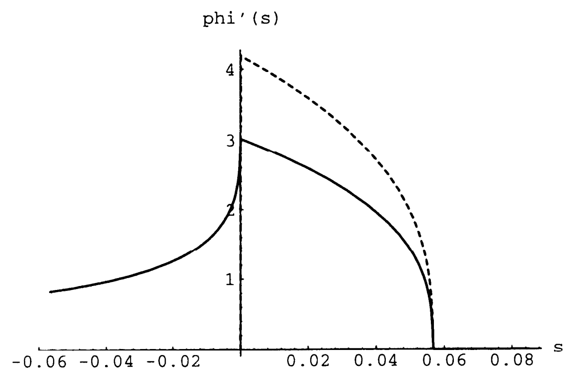

Figure 4: Derivative of underlying Brownian Motion for FBM (solid line) and Levy's version of FBM (dashed line) when $H = 0.85$ and $\mu = 100$.

level x at time τ, as:

$$\phi_t^L(s) = \begin{cases} \dfrac{2Hx}{\tau^{2H}} \displaystyle\int_0^s (s-y)^\alpha (\tau - y)^\alpha dy & s \le \tau \\ \dfrac{2Hx}{\tau^{2H}} \displaystyle\int_0^\tau (s-y)^\alpha (\tau - y)^\alpha dy & s > \tau \end{cases} \tag{29}$$

Simple calculations show that the asymptotics for the Levy version of FBM have the same expression as in (8) and (17) and therefore the buffer asymptotics are still given by (24). Figure 1, 2 and 3 show with dashed lines the behaviour of:

$$Q_b(s) = \max\left\{0, \phi_b^L(s) - \mu s\right\}, \tag{30}$$

for different values of H and μ. The essential difference between the two models is that the model with stationary increments has a 'history' of infinite past which it can rig in order to maximise its chances of reaching a high level; in some sense, it has a 'headstart' over the Levy version, which starts at time zero with no past. The overflow event is thus 'smoother' for the model with stationary increments.

8 Conclusions

We have analysed a scaled version of fractional Brownian Motion and determined the most likely way in which it reaches a given level and the most likely path it will follow thereafter given that it has. This was made possible by a representation of the scaled FBM as an additive functional of an underlying *moderately scaled* BM; the large deviations results for the FBM could then be deduced by contraction from a moderate deviations analogue of Schïlder's theorem. We then applied this to the corresponding queueing model to determine the most likely path to overflow (in a large buffer) and the most likely path followed after overflow has occurred. We drew the following conclusions about the *shape* of this path. It depends only on the Hurst parameter: if H is close to $1/2$ the paths followed are almost linear (as one would expect, since the case $H = 1/2$ is BM, where the most likely trajectories are indeed linear); as H increases, the path to overflow becomes more and more 'S-shaped', and the system is slower to recover after an overflow event.

References

[1] 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*, Vol. 2, No. 1, pp. 1-15, February 1994.

[2] B. Mandelbrot and J. Van Ness, "Fractional Brownian Motions, Fractional Gaussian Noises and Applications", *SIAM Review*, Vol. 10, No. 4, pp. 422-437, October 1968.

[3] G. Samorodnitsky, M. S. Taqqu, *Stable Non-Gaussian Random Processes*, Chapman & Hall, 1994.

[4] A. Dembo, O. Zeitouni, *Large Deviations Techniques and Applications*, Jones and Bartlett, 1992.

[5] Jean-Dominique Deuschel and Daniel W. Stroock. *Large deviations.* Academic Press, 1989.

[6] N. G. Duffield, N. O'Connell, "Large Deviations and Overflow Probability for the General Single-Server Queue With Applications", *Math. Proc. Cambridge Philos. Society,* , Vol. 118, No. 2, pp. 363-374, 1995.

[7] G. F. Simmons, *Differential Equations With Applications and Historical Notes*, second edition, McGrow-Hill, 1991.

[8] I. Norros, "On the Use of Fractional Brownian Motion in the Theory of Connectionless Networks", *IEEE Journal on Selected Areas in Communications*, Vol. 13, No. 6, pp. 953-962, August 1995.

[9] C. S. Chang, D.D. Yao, T. Zajic, "Large Deviations, Moderate Deviations, and Queues with Long-Range Dependent Input", to appear in *Journal of Applied Probability.*