

Link Dimensioning for Fractional Brownian Input

Jiongze Chen
PhD student, Electronic
Engineering Department, City
University of Hong Kong,
Hong Kong SAR
Email: cjz12110@gmail.com
Telephone: (852) 98084853

Ronald G. Addie
Department of Mathematics
and Computing, University of
Southern Queensland, Australia
Email: addie@usq.edu.au
Telephone: (61) 7-4631-5520
Fax: (61) 7-4631-5550

Moshe Zukerman
Electronic Engineering Department,
City University of Hong Kong,
Hong Kong SAR
Email: m.zu@cityu.edu.hk
Telephone: (852) 3442-4243
Fax: (852) 2788-7791

Abstract—This paper provides a new formula for dimensioning of a link fed by fractional Brownian input. This formula is obtained based on another new approximate result for the stationary workload distribution of a queue loaded by fractional Brownian input. An efficient approach to simulate such a queue is also presented. Agreement between the analytical and the simulation results has been demonstrated numerically.

Index Terms—fractional Brownian noise, fractional Brownian motion, long range dependence, quality of service (QoS), link capacity dimensioning.

I. INTRODUCTION

It has been well established that Internet traffic exhibits long range dependence (LRD) characteristics [1]–[5]. Furthermore, core and metropolitan Internet links are shared by a large number of users, so by the central limit theorem, the traffic on such links which represents multiplexing of traffic generated by many users can be assumed to follow a Gaussian process for the purpose of performance evaluation and link capacity dimensioning [6]. In addition, Internet traffic is transported based on the store and forward principle, where packets are stored in router buffers before they are forwarded towards their destination. Therefore, a queue fed by fractional Brownian motion (fBm) input has been considered a fundamentally important model for Internet queueing performance analysis and capacity assignment and has attracted significant attention [7]–[15]. To date, despite the effort, only asymptotic results are available for queueing performance of fBm queues.

This paper has two distinguishable parts. The first part comprises sections II, III and IV, where we consider the theoretical single server queue fed by fBm traffic. In Section II, a new approximation for the stationary workload distribution of an fBm queue is provided. In Section III, we describe an efficient method to simulate an fBm queue. Numerical results are presented in Section IV to demonstrate the accuracy of our new fBm queue workload distribution formula.

In the second part, sections V and VI, we aim to provide a first step towards bridging the gap between theory and practice, where we apply the theoretical results of the first part to link dimensioning and explain the limitations of the

theory. In particular, in Section V we use our new queueing performance approximation to derive an elegant new formula for link dimensioning that gives the capacity required to meet QoS measured as the proportion of time that the workload distribution is above a given threshold. Then in Section VI we discuss a limitation of the practical application of the model due to the presence of significant quantities of negative traffic in the theoretical model in some situations.

II. A NEW ANALYTICAL RESULT OF AN FBM QUEUE

We consider a single server queue with an infinite buffer fed by an fBm input process with Hurst parameter H , variance σ_1^2 and drift λ . Specifically, the variables σ_1^2 and λ are the variance and the mean of the amount of work arriving during a time interval of length 1. The service rate, denoted by τ , is assumed constant. Let μ be the mean net input during a time interval of length 1, i.e., $\mu = \lambda - \tau$. For stability we assume $\mu < 0$. Henceforth, a time interval of length 1 will be called 1 second [sec.]. Let Q be the steady state queue size.

For the case $H = 0.5$, the complementary distribution function of Q , namely, the probability that the fBm queue size exceeds x , denoted $P(Q > x)$, is well known and is given by [16]:

$$P(Q > x) = \exp \frac{2\mu}{\sigma_1^2} x, \quad x \geq 0. \quad (1)$$

For $H \neq 0.5$, despite considerable efforts, there is no exact result for $P(Q > x)$.

Of the known results, a first approximation/bound for the complementary distribution function is from [9]:

$$P(Q > x) \approx \exp \left(- \frac{x^{2-2H}(1-H)^{2H-2}|\mu|^{2H}}{2H^2H\sigma_1^2} \right), \quad x \geq 0. \quad (2)$$

The authors of [8] showed that this approximation holds in the sense that

$$\lim_{x \rightarrow \infty} \frac{1}{x} \log (\text{LHS/RHS}) = 0. \quad (3)$$

A second approximation results from from [14, Theorem 1, Equation (9)] with $\alpha = 2H$, $\beta = 1$, together with the asymptotic approximation for the tail of a Normal distribution,

$$1 - \Phi(x) \sim x^{-1} \exp \left(- \frac{x^2}{2} \right),$$

giving

$$P(Q > x) \approx$$

$$\tilde{C} x^{\frac{2H^2-3H+1}{H}} e^{-\frac{x^{2-2H}(1-H)^{2H-2}|\mu|^{2H}}{2H^{2H}\sigma_1^2}}, \quad x \geq 0, \quad (4)$$

where \tilde{C} is a certain constant. This approximation holds in the sense that

$$\lim_{x \rightarrow \infty} \text{LHS/RHS} = 1. \quad (5)$$

The authors of [14] did not provide a way to compute \tilde{C} . Therefore we estimate this constant by numerically fitting (4) to simulation results.

In this paper we propose the following approximation, which is expressed in terms of the density rather than the complementary distribution function. Let

$$P(Q > u) \approx \int_u^\infty cf(x)dx. \quad (6)$$

where

$$f(x) = x^{\frac{1-2H}{H}} e^{-\frac{x^{2-2H}(1-H)^{2H-2}|\mu|^{2H}}{2H^{2H}\sigma_1^2}}, \quad x \geq 0, \quad (7)$$

and

$$\begin{aligned} c^{-1} &= \int_0^\infty f(x)dx \\ &= \int_0^\infty x^{\beta-1} e^{-\alpha x^\nu} dx \end{aligned}$$

in which

$$\begin{aligned} \alpha &= \frac{(1-H)^{2H-2}|\mu|^{2H}}{2H^{2H}\sigma_1^2}, \\ \nu &= 2 - 2H \end{aligned}$$

and

$$\beta = \frac{1-H}{H};$$

so if we set

$$y = x^\nu,$$

hence

$$x = y^{\frac{1}{\nu}},$$

and

$$\begin{aligned} dx &= \frac{1}{\nu} y^{\frac{1}{\nu}-1} dy, \\ c^{-1} &= \int_0^\infty y^{\frac{\beta-1}{\nu}} e^{-\alpha y^{\frac{1}{\nu}}} \frac{1}{\nu} y^{\frac{1}{\nu}-1} dy \\ &= \frac{1}{\nu} \int_0^\infty y^{\frac{\beta}{\nu}-1} e^{-\alpha y} dy \\ &= \nu^{-1} \alpha^{-\frac{\beta}{\nu}} \Gamma\left(\frac{\beta}{\nu}\right), \end{aligned}$$

in which

$$\Gamma(z) = k^z \int_0^\infty t^{z-1} e^{-kt} dt.$$

The approximation (6) holds in the sense that $\lim_{x \rightarrow \infty} \text{LHS/RHS} = B$ for some constant B . This follows from (5) because the RHS of (7) is the derivative of the RHS of (4), up to a constant multiple. Although the

constant B here may not be equal to 1, which means that (6) may be less accurate than (4) for large x , the new formula has the advantage that it is a distribution. In particular, whereas (4) cannot be accurate for small x unless $H = 0.5$, it is feasible that (6) is accurate for small x . The simulation results appear to confirm that this is the case. Another benefit of (6) over that of (4) is that we provide a way to derive c while (4) does not provide a way to derive \tilde{C} . In section IV below we compute \tilde{C} by fitting (4) to simulation results for large x .

III. SIMULATION OF AN FBM QUEUE

The challenge in simulating an fBm queue is how to adapt the continuous time concept of fBm to the discrete-time implementation of a computer simulation that considers the queue size at the endpoints of consecutive intervals each of size Δt [sec.]. This forces us to consider a limit, namely, the queueing performance limit as we perform a sequence of such simulations where Δt becomes smaller and smaller.

A simulation of an fBm queue requires the generation of a long fractional Brownian noise (fBn) sequence which is time consuming. There is a significant body of work on how to generate fBm traffic [17]–[21]. We do not intend to contribute in this area here. Instead we use an fBn sequence obtained by the Hosking recursive method using the code of Dieker [22]. For computational efficiency, we do not generate a new fBn sequence for each simulation, but instead use the same sequence. For each scenario (given H value), we use an independent fBn sequence of length 2^{22} seconds. The fBn sequence is characterized by its Hurst parameter H and its variance v . As we reduce Δt we still use the same sequence, which implies that we must have the same H and v . Clearly, the variance of the work arrived is reduced as Δt is reduced, so to maintain the same numerical value of v for the variance, we need to scale down the “units” of work. As a consequence, in every simulation we will *change* some simulation parameters so that each remains consistent with the *same* continuous time model.

As is often the case in queueing simulations, the basic algorithm of our discrete-time queue simulation is *Lindley's equation*,

$$Q_{n+1} = \max(0, Q_n + U_n + m), \quad (8)$$

in which Q_n denotes the queue length at the end of the n th time-interval and $m = m(\Delta t)$ is the net mean input per time-interval. The variable U_n denotes the difference between the value of the fBm at the end of the time-interval n and the value of the fBm at the beginning of the time-interval n . U_n can be negative.

The problem tackled in this section is as follows. Given that the vector U contains a fractional Brownian noise process with Hurst parameter H , and variance v , what value of m should be used, and what *scaling* of Q_n should be adopted to interpret equation (8) as an fBm queue in which the time interval is Δt [sec.], the net mean input is μ units of work per one [sec.], and the variance in one [sec.] is σ_1^2 ?

To adapt (8) to the continuous time model, fBm, we need to envisage a limit as we alter our interpretation of the discrete time of our simulation to correspond to shorter and shorter time intervals. Since the time consuming aspect of the simulations is the generation of fractional Brownian motion (or noise – the two are related by differencing or cumulative integration), we shall assume that a discrete-time fractional Brownian motion process has been generated *in advance* and our plan is to re-use this fixed dataset for multiple purposes of the following sort:

- (i) more and more accurate simulation of an fBm queue, in the sense that the loss of accuracy due to discretising time is reduced successively;
- (ii) simulation of queues with different variance or net mean input.

To clarify, v is the variance per unit time of the fBm sequence and σ_1^2 is the variance of the amount of work arriving within one second in the simulated model. In our first simulation, $\Delta t = 1$ [sec.] so $v = \sigma_1^2$. In the second $\Delta t = 0.1$ [sec.], etc. Again, we always use the same original fBn sequence, but in each simulation we change m , and we interpret work using a different scale factor.

A. Choice of m

According to the Hurst formula, the variance in an interval of length Δt should be $\sigma_1^2(\Delta t)^{2H}$, but actually it is v , so we must interpret each nominal unit of work as representing $\sigma_1(\Delta t)^H/\sqrt{v}$. Therefore we should set

$$m = \frac{\mu\Delta t\sqrt{v}}{\sigma_1(\Delta t)^H}.$$

B. Interpretation of Q

We now completed the simulation using a new m and obtained the function $P(Q > x)$ based on that new m . But this is not the correct function because the work arrives has also to be scaled. To do this we need to divide x by a scaling factor S .

As discussed above, the units in the simulated system actually correspond to $\sigma_1(\Delta t)^H/\sqrt{v}$ units of work in the model we are trying to simulate, so we should divide the x values by

$$S = \frac{\sigma_1(\Delta t)^H}{\sqrt{v}}$$

in order to be able to interpret them as applicable to the intended fBm model. Then we consider the function $P(Q > x/S)$ instead of the function $P(Q > x)$.

IV. VALIDATION OF THE WORKLOAD DISTRIBUTION

In this section we present numerical results that demonstrate the accuracy of our simulation and of our approximation. All simulation results are provided with 95% confidence intervals based on the Student-t distribution.

A. Validation of the simulation

To validate the simulation, in Figure 1 using the known exact result (1) for the case $H = 0.5$, we demonstrate how the sequence of simulations as described above approach the exact result using only one fBn sequence. We note that in the case $\Delta t = 0.01$ the analytical results are within the simulation confidence intervals.

B. Validation of the analytical formula

In Figures 2 – 5, we present results for $H = 0.3, 0.4, 0.6, \text{ and } 0.7$, respectively, that demonstrate the accuracy and robustness of our approximation. In most instances the approximation is within the confidence intervals of the simulation over the entire range and in all instances the discrepancies which do exist are quite small. We also demonstrate that, as expected, the existing asymptotes are not accurate for the full range of parameters in all cases. Figures 6 and 7 present clearer results for $H = 0.3$ and 0.4 with $x > 1$. Results for $x \geq 4$ are not presented due to limitation of simulation accuracy. Notice the increase length of the confidence intervals as we approach $x = 4$. Figure 8 that our approximation is quite significantly more accurate than the previous results for $H = 0.7$ when x is small.

C. Discussion of results

First of all it should be observed that the formula presented here, based on comparison with simulation, appears to be more accurate than the existing alternative formulae in all the situations which have been tested, up to now.

Secondly, the new formula has good accuracy for the full range of parameters.

Thirdly, although it might appear that the Husler-Piterburg asymptote [14] also provides adequate accuracy over the full practical range of the parameters (for our purposes, for example, its weakness when H and x are small is not practically important), this formula cannot readily be used independently of simulations because it makes use of a constant for which there is no clear method of calculation. In order to use the Husler-Piterburg asymptote in the above comparisons, this constant was estimated by simulations as discussed in Section II.

V. LINK DIMENSIONING

Having validated the analytical formula we have derived for the workload distribution of a queue fed by fBm traffic, we are now in position to evaluate the capacity required C to serve a link fed by fBm traffic such that the probability $\varepsilon = P(Q > q)$ for a given queue threshold q is below a given margin. Hence, we are interested in the link capacity required C to meet our QoS measures ε and q .

A. A theoretical formula for link dimensioning

Our formula is:

$$P(Q > q) \approx \int_q^\infty cf(x)dx.$$




Fig. 1. Overflow probabilities based on exact formula vs. Simulation results for $H = 0.5, \sigma_1^2 = 1, \mu = -0.5$ with $\Delta t = 1, 0.1$ and 0.01 .




Fig. 2. Overflow probabilities based on the asymptotes of [9] and [14], and our approximation vs. simulation results for $H = 0.3, \sigma_1^2 = 1, \mu = -0.5$.




Fig. 3. Overflow probabilities based on the asymptotes of [9] and [14], and our approximation vs. simulation results for $H = 0.4, \sigma_1^2 = 1, \mu = -0.5$.




Fig. 4. Overflow probabilities based on the asymptotes of [9] and [14], and our approximation vs. simulation results for $H = 0.6, \sigma_1^2 = 1, \mu = -0.5$.




Fig. 5. Overflow probabilities based on the asymptotes of [9] and [14], and our approximation vs. simulation results for $H = 0.7, \sigma_1^2 = 1, \mu = -0.5$.

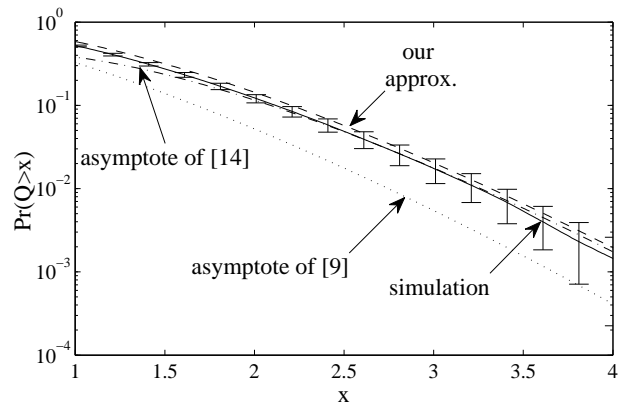


Fig. 6. Overflow probabilities based on the asymptotes of [9] and [14], and our approximation vs. simulation results when x is large for $H = 0.3, \sigma_1^2 = 1, \mu = -0.5$.

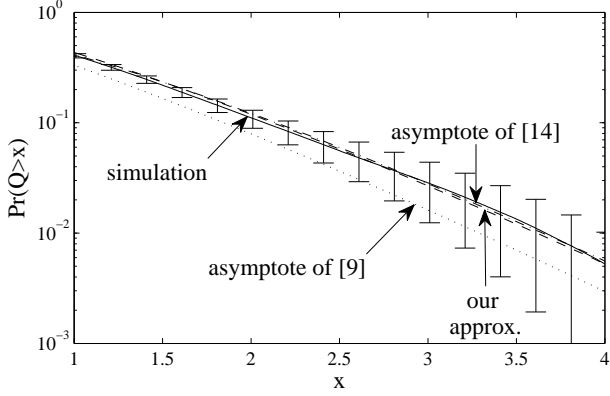


Fig. 7. Overflow probabilities based on the asymptotes of [9] and [14], and our approximation vs. simulation results when x is large for $H = 0.4, \sigma_1^2 = 1, \mu = -0.5$.

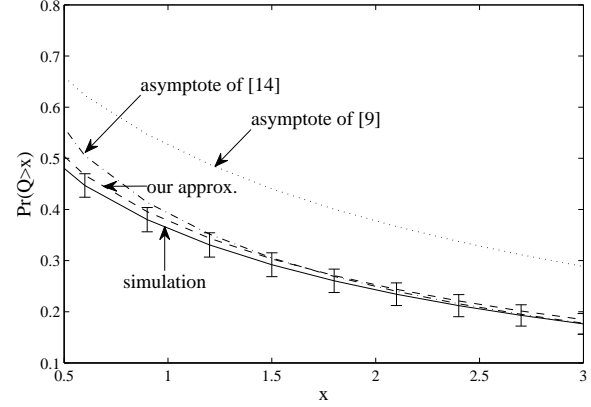


Fig. 8. Overflow probabilities based on the asymptotes of [9] and [14], and our approximation vs. simulation results when x is small for $H = 0.7, \sigma_1^2 = 1, \mu = -0.5$.

Let us define the regularised incomplete Gamma function by

$$G(a, x) = \left(\int_x^\infty t^{a-1} e^{-t} dt \right) / \Gamma(a). \quad (9)$$

By substituting kt for t , we find

$$\Gamma(a)G(a, x) = k^a \int_{x/k}^\infty t^{a-1} e^{-kt} dt.$$

Then we can derive a dimensioning formula from (6) as follows:

$$\begin{aligned} P(Q > q) &\approx \int_q^\infty cf(x)dx \\ &= \frac{\nu\alpha^{\frac{\beta}{\nu}}}{\Gamma\left(\frac{\beta}{\nu}\right)} \int_q^\infty y^{\beta-1} e^{-\alpha y^\nu} dy \\ &\text{which, using the substitution } u = y^\nu, \text{ in which case } du = \nu y^{\nu-1} dy, \text{ or putting it another way, } du = \nu u^{\frac{\nu-1}{\nu}} dy \\ &= \frac{\alpha^{\frac{\beta}{\nu}}}{\Gamma\left(\frac{\beta}{\nu}\right)} \int_{q^\nu}^\infty u^{\frac{\beta-1}{\nu}} u^{\frac{1-\nu}{\nu}} e^{-\alpha u} du \\ &= \frac{\alpha^{\frac{\beta}{\nu}}}{\Gamma\left(\frac{\beta}{\nu}\right)} \int_{q^\nu}^\infty u^{\frac{\beta}{\nu}-1} e^{-\alpha u} du \\ &= G\left(\frac{1}{2H}, \alpha q^\nu\right) \end{aligned} \quad (10)$$

Let us define the inverse regularised incomplete Gamma function, $G^{-1}(\alpha, y)$, by the property

$$G^{-1}(\alpha, G(\alpha, x)) = G(\alpha, G^{-1}(\alpha, x)) = x. \quad (11)$$

Now since we have defined $\varepsilon = P(Q > q)$, and recall that $\nu = 2 - 2H$, we obtain:

$$\alpha \approx G^{-1}\left(\frac{1}{2H}, \varepsilon\right) q^{2H-2}. \quad (12)$$

Let us now use the original defining equation for α , which expresses it in terms of μ , to determine a formula for μ and

hence then C :

$$\begin{aligned} \alpha &= \frac{(1-H)^{2H-2} |\mu|^{2H}}{2H^{2H} \sigma_1^2} \\ \implies (1-H)^{2H-2} |\mu|^{2H} &= 2\alpha H^{2H} \sigma_1^2 \\ \implies |\mu| &= \left(\frac{2\alpha H^{2H} \sigma_1^2}{(1-H)^{2H-2}} \right)^{\frac{1}{2H}}. \end{aligned} \quad (13)$$

Now since $\mu = \lambda - C$, where C is the capacity and λ denotes denotes the mean rate of the traffic,

$$\begin{aligned} C - \lambda &= \left(\frac{2\alpha H^{2H} \sigma_1^2}{(1-H)^{2H-2}} \right)^{\frac{1}{2H}} \\ \text{so} \\ C &= \lambda + \left(\frac{2H^{2H} \sigma_1^2 G^{-1}((2H)^{-1}, \varepsilon)}{(1-H)^{2H-2}} \right)^{\frac{1}{2H}} q^{-\frac{1-H}{H}}. \end{aligned} \quad (14)$$

We have obtained a simple and elegant formula for link capacity loaded by fBm traffic. From this formula simple relationships can be observed. One observation is that the required spare capacity beyond the arrival rate (λ) is independent of λ . Another simple observation is that the spare capacity relates to σ_1 according to $\sigma_1^{1/H}$. In next subsection we provide numerical results to illustrate such relationships.

B. Numerical results

Numerical results of link dimensioning for a range of examples using the method developed above are presented in Figures 9–13. The results presented are for the total capacity C , but it is important to keep in mind the concept of spare capacity $C - \lambda$. In all the figures except Fig. 12 we set $\lambda = 1$. This could represent data rate in order of Gb/s, e.g., one OC-192 or OC-768 with rate of 10 Gb/s or 40 Gb/s. Then $q = 0.1$ would represent QoS measure of 100 ms.

In Figure 9 we illustrate the total capacity C required as a function of the second QoS measure ε within the range $10^{-5} - 1$. As expected the spare capacity required, i.e. $C - 1$, is




Fig. 9. Capacity versus ε by (14) for $H = 0.8$, $\sigma_1 = 0.05$, $\lambda = 1$ and $q = 0.1$.




Fig. 10. Capacity versus H by (14) for $\sigma_1 = 0.05$, $\lambda = 1$, $q = 0.1$ and $\varepsilon = 0.001$.

reduced with the QoS measure ε . For the case $\lambda = 1$, the spare capacity $C - 1$ approaches zero as the QoS is further and further relaxed.

In Figure 10 we illustrate that the total capacity C increases with the Hurst parameter H . This is also expected as stronger correlation in the traffic stream means higher queueing delay.

In Figure 11 we illustrate that the total capacity C increases with the standard deviation parameter σ_1 . This is also expected as stronger variation in the traffic stream, normally, means higher queueing delay.

As observed in (14), total capacity C increases linearly with λ . This is illustrated in Figure 12.

In Figure 13 we illustrate that the total capacity C decreases as the QoS parameter q increases. Namely, allowing more delay and more buffering relaxed delay requirement, so less capacity is expected to be required. The space capacity required approached zero as queue is allowed to approach infinity. This is consistent with known results in elementary Markovian queues.

VI. DISCUSSION

The fractional Brownian motion model is not universally appropriate to Internet traffic. An important weakness of this




Fig. 11. Capacity versus σ_1 by (14) for $H = 0.8$, $\lambda = 1$, $q = 0.1$ and $\varepsilon = 0.001$.




Fig. 12. Capacity versus λ by (14) for $H = 0.8$, $\sigma_1 = 0.05$, $q = 0.1$ and $\varepsilon = 0.001$.




Fig. 13. Capacity versus q by (14) for $H = 0.8$, $\sigma_1 = 0.05$, $\lambda = 1$ and $\varepsilon = 0.001$.

model is that for certain ranges of the parameters it exhibits the phenomenon of *negative traffic* to an excess. A small amount of negative traffic will always be present in the fBm model, but since real networks have no negative traffic at all, it is inappropriate to use a model which has large amounts of it. When σ_1 is large relative to λ , or when $H \ll 0.5$, and especially when both of these conditions hold, the model will exhibit large amounts of negative traffic when observed at small time scales.

It may be possible to re-interpret an fBm model in a manner which alleviates this problem and enables us to use this model for a wider range of parameters. However, in the absence of such a strategy it will be necessary to confine the use of the fBm model to situations where $H \geq 0.6$, $q > 0.1$, and $\sigma_1 \leq 0.05$. The constraint on H is quite reasonable, since H values smaller than this are rarely if ever observed in real networks. The constraint on σ_1 means that we must confine applications to core links of networks, where the traffic is relatively smooth. This was indeed, always our intention. The constraint on q means that we are assuming buffers large enough to hold the traffic arriving in 0.1 seconds. The model can be used with smaller buffers than this if the traffic is smoother, i.e. σ_1 is less.

VII. CONCLUSION

We have considered a queue fed by fBm input and derived new results for queueing performance and link dimensioning. We have also described an efficient approach to simulate such a queue. Agreement between the analytical and the simulation results have been demonstrated and a numerical comparison with existing asymptotes has been presented. We have also presented numerical results for a range of examples for link dimensioning based on our queueing analysis.

Finally, we would like to comment that the work here is only an initial step to bridge the gap between theory and practice. Much work is still required to understand the limitations and the usefulness of the fBm model in network design. Such understanding must be established using more realistic models and actual measurements.

ACKNOWLEDGMENT

The authors would like to thank Ilkka Norros and Ward Whitt for advice on relevant literature.

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 Trans. Networking*, vol. 2, no. 1, pp. 1–15, Feb. 1994.
- [2] R. G. Addie, M. Zukerman, and T. Neame, "Fractal traffic: measurements, modelling and performance evaluation," in *Proc. IEEE INFOCOM '95*, vol. 3, Apr. 1995, pp. 977–984.
- [3] J. Beran, R. Sherman, M. S. Taqqu, and W. Willinger, "Long-range dependence in variable-bit-rate video traffic," *IEEE Trans. Commun.*, vol. 43, no. 2, 3, 4, pp. 1566–1579, Feb./March/April 1995.
- [4] V. Paxson and S. Floyd, "Wide-area traffic: the failure of Poisson modelling," *IEEE/ACM Trans. Networking*, vol. 3, no. 3, pp. 226–244, Jun. 1995.
- [5] M. E. Crovella and A. Bestavros, "Self-similarity in world wide web traffic: evidence and possible causes," *IEEE/ACM Trans. Networking*, vol. 5, no. 6, pp. 835–846, Dec. 1997.

- [6] M. Zukerman, T. D. Neame, and R. G. Addie, "Internet traffic modeling and future technology implications," in *Proc. IEEE INFOCOM 2003*, vol. 1, Apr. 2003, pp. 587–596.
- [7] R. Dijkerman and R. Mazumdar, "On the correlation structure of the wavelet coefficients of fractional Brownian motion," *IEEE Trans. Inf. Theory*, vol. 40, no. 5, pp. 1609–1612, Sep. 1994.
- [8] N. G. Duffield and N. O'Connell, "Large deviations and overflow probabilities for the general single-server queue, with applications," *Proceedings of the Cambridge Philosophical Society*, vol. 118, no. 2, pp. 363–374, 1995.
- [9] I. Norros, "A storage model with self-similar input," *Queueing Syst.*, vol. 16, no. 3–4, pp. 387–396, Sep. 1994.
- [10] —, "On the use of fractional Brownian motion in the theory of connectionless networks," *IEEE J. Sel. Areas Commun.*, vol. 13, no. 6, pp. 953–962, Aug. 1995.
- [11] —, "Four approaches to the fractional Brownian storage," *Fractals in Engineering*, pp. 154–169, 1997.
- [12] O. Narayan, "Exact asymptotic queue length distribution for fractional Brownian traffic," in *Adv. Perform. Anal.*, vol. 1, no. 1, 1998, pp. 39–63.
- [13] L. Massoulié and A. Simonian, "Large buffer asymptotics for the queue with fractional Brownian input," *J. Appl. Probab.*, vol. 36, no. 3, pp. 894–906, Sep. 1999.
- [14] J. Hüslér and V. Piterbarg, "Extremes of a certain class of Gaussian processes," *Stochastic Processes and their Applications*, vol. 83, no. 2, pp. 257 – 271, Oct. 1999.
- [15] Y. Jin, S. Bali, T. Duncan, and V. Frost, "Predicting properties of congestion events for a queueing system with fBm traffic," *IEEE/ACM Trans. Networking*, vol. 15, no. 5, pp. 1098 –1108, Oct. 2007.
- [16] J. M. Harrison, *Brownian motion and stochastic flow systems*. New York: John Wiley and Sons, 1985.
- [17] B. B. Mandelbrot and J. W. V. Ness, "Fractional brownian motions, fractional noises and applications," *Society for Industrial and Applied Mathematics*, vol. 10, no. 4, pp. 422–437, Oct 1968.
- [18] J. Hosking, "Modeling persistence in hydrological time series using fractional differencing," *Water resources research*, vol. 20, no. 12, pp. 1898–1908, Aug. 1984.
- [19] C. Dietrich and G. Newsam, "Fast and exact simulation of stationary Gaussian processes through circulant embedding of the covariance matrix," *SIAM Journal Sci. Comput.*, vol. 18, no. 4, pp. 1088–1107, Jul. 1997.
- [20] V. Paxson, "Fast, approximate synthesis of fractional gaussian noise for generating self-similar network traffic," *ACM Sigcomm Comput. Commun. Rev.*, vol. 27, no. 5, pp. 5–18, Oct. 1997.
- [21] W. Whitt, *Stochastic-process limits*. New York, USA: Springer-Verlag, 2002, pp. 124–130.
- [22] A. B. Dieker, "Fractional Brownian motion." [Online]. Available: <http://www2.isye.gatech.edu/~adieker3/fbm/index.html>