Synthesis of Fractional Gaussian Noise Using Linear Approximation for Generating Self-Similar Network Traffic*

SERGIO LEDESMA[1] AND DERONG LIU[2]
1. Department of Electrical and Computer Engineering
Stevens Institute of Technology, Hoboken, NJ 07030
Email: sledesma@stevens-tech.edu
2. Department of Electrical Engineering and Computer Science
University of Illinois, Chicago, IL 60607
liu@eecs.uic.edu

Abstract

The present paper focuses on self-similar network traffic generation. Network traffic modeling studies the generation of synthetic sequences. The generated sequences must have similar features to the measured traffic. Exact methods for generating self-similar sequences are not appropriate for long traces. Our main objective in the present paper is to improve the efficiency of Paxson's method for synthesizing self-similar network traffic. Paxson's method uses a fast, approximate synthesis for the power spectrum of the FGN and uses the inverse Fourier transform to obtain the time-domain sequences. We demonstrate that a linear approximation can be used to determine the power spectrum of the FGN. This linear approximation reduces the complexity of the computation without compromising the accuracy in synthesizing the power spectrum of the FGN. Our results show that long traces can be generated in much less time. To compare our method with existing ones, we will measure the running time in generating long and short sample paths from the FGN. We will also conduct experiments to show that our method can generate self-similar traffic for specified Hurst parameters with high accuracy.

*This work was supported in part by the National Science Foundation under Grant ECS-9996428 and in part by Blonder Broadcasting Corporation, Morganville, New Jersey.

1 Introduction

Data traffic is the main component of modern computer communication systems, and traffic models are of crucial importance for assessing their performance [19]. In practice, stochastic models of traffic streams are relevant to network traffic engineering and performance analysis, and they are widely used in predicting system performance. The basic systems, of which traffic is a major ingredient, are queueing systems. Traditional traffic models have often been devised and chosen for the analytical tractability they induce on the corresponding queueing systems. While originally the validity and efficacy of models for modern high-speed network traffic were difficult to establish due to the unavailability of empirical data, very large sets of traffic

measurements from working networks have become available.

Recently, the notion of self-similarity has been shown to apply to a variety of traffic including wide area network (WAN) traffic [16], local area network (LAN) traffic [13], [21], frame data generated by variable-bit-rate (VBR) video encoders [4], [10], and the World Wide Web (WWW) transfers [5]. We briefly review in the following some applications where self-similar traffic models have been used.

- WAN traffic. When modeling network traffic, packet and connection arrivals are often assumed to be Poisson processes in the past. Some recent studies have shown, however, that the distribution of packet interarrivals clearly differs from the exponential distribution. In [16], it is concluded that Poisson arrival processes are quite limited in modeling the burstiness of data traffic, especially when many sources are multiplexed together. It is shown that the WAN traffic is much burstier than Poisson model's prediction over all time scales. The greater burstiness of data traffic has implications for many aspects of congestion control and network performance. It is suggested in [16] that the burstiness of empirical traffic meshes well with self-similar network traffic models.

- LAN traffic. An analysis of Ethernet LAN traffic is performed in [13] and [21]. It is shown that irrespective of when and where the Ethernet measurements are collected, the LAN traffic is self-similar with different degrees of self-similarity depending on the load on the network. Traffic models based on self-similar stochastic processes are presented. It is also shown in the papers that the burstiness of LAN traffic typically intensifies as the number of active traffic sources increases, contrary to commonly held views.

- VBR video traffic. An analysis of a few millions of encoded video frame data generated by VBR video encoders shows that the VBR video traffic appears to be statistically self-similar [4], [10]. One advantage of the modeling approach for video traffic based on long-range dependence/self-similarity is that it removes entirely the subjective identification of scenes and scene changes in some currently available VBR video traffic modeling literature.

- WWW transfers. In [5], it is observed that traffic due to the WWW transfers shows characteristics that are consistent with self-similarity. They trace the genesis of Web traffic self-similarity to the heavy-tailed distribution of available file sizes in the Web. It has been argued in [5] and others that transferring files whose sizes are drawn from a heavy-tailed distribution will generate self-similarity in network traffic.

All the above studies provide inspirations for an investigation of a new type of traffic models. In particular, it is encouraged to use stochastic processes which possess scaling properties and which have distributions with infinite moments. Self-similar structures, and hence fractal processes, have been employed for this purpose. The study of [21] and others have concerned about whether it is possible to statistically distinguish between measured network traffic and model generated traffic. Actual traffic exhibits correlations over a wide time span (i.e., *long-range*

dependence), while traditional traffic models typically focus on a very limited range of time spans and are thus *short-range dependent* in nature.

A fundamental feature of a synthetic traffic generator is to provide reliable data that can be used instead of experimental traffic traces to assess the performance and/or to dimension the network elements that deal with such traffic. Another important issue for the success of traffic modeling in teletraffic practice is the ability to quickly generate traces of synthetic traffic from a chosen traffic model. While exact methods for generating traces from some basic self-similar processes exist, they are in general only appropriate for short traces and becomes impractical when the required number of samples becomes exceedingly large.

This paper presents a *fast* method for generating long traces of self-similar traffic. Our method will also guarantee to generate self-similar traffic with high degree of accuracy in terms of burstiness measures (i.e., the Hurst parameter estimated from the traffic trace). There are three mathematical models used often to model the self-similarity effect: The fractional Gaussian noise (FGN), the fractional Brownian motion (FBM), and the fractional autoregressive integrated moving average (F-ARIMA) process. We will investigate in this paper the use of FGN for generating self-similar network traffic. Our goal is to develop a fast method in generating the spectral density function for the FGN so that the FGN can be generated by taking the inverse Fourier transform (FFT).

Approaches for generating self-similar sequences based on the three mathematical models mentioned above have been reported in the literature. Paxson [15] reported a fast approach using FFT to generate the FGN. The present approach constitutes an improvement to Paxson's and our simulation results in this paper will show that the present approach and Paxson's approach can generate high quality self-similar sequences (i.e., the Hurst parameters of the generated sequences agree with the target values). An approach for generating self-similar traces based on the FBM was reported in [12]. However, our simulation results show that the approach of [12] fails to provide stationary increments which violates one of the desired properties of FBM [11]. Hosking's method reported in [9] for generating self-similar sequences based on the F-ARIMA is known to be very slow. In addition, due to the iterative use of equations in this approach, an improvement to the speed of computation seems difficult. All these motivate the choice of FGN to model the self-similar processes in the present study.

The present paper is organized as follows. In Section 2, we present a brief review of the FGN. In Section 3, we develop a method using a linear approximation to calculate the power spectrum of the FGN. In Section 4, we compare the method developed herein with existing methods for generating self-similar traffic in terms of the computation time. We will also show the accuracy of the present approach in generating self-similar traffic with given Hurst parameters. Finally, in Section 5, we conclude the present paper with a few pertinent remarks.

2 The Fractional Gaussian Noise

Long-range dependence in a time series is the presence of a significant correlation between observations of signals separated by large time spans. It is closely linked with self-similar stochastic processes and random fractals which have been considered exten-

sively, though only recently, for signal processing applications.

Let $X = \{X_t : \ t = 0, 1, 2, \cdots\}$ be a covariance stationary (also called wide-sense stationary) stochastic process with mean μ, variance σ^2 and autocorrelation function $r(k)$, $k = 0, 1, 2, \cdots$. In particular, we assume that X has an autocorrelation function of the form

$$r(k) \sim k^{-(2-2H)} L(k) \quad \text{as} \quad k \to \infty \quad (1)$$

where H is called the *Hurst parameter* and $L(\cdot)$ is slowly varying at infinity, that is,

$$\lim_{t \to \infty} L(zt)/L(t) = 1 \ \text{ for all } \ z > 0 \quad (2)$$

An example of such slowly varying functions which satisfies (2) is given by $L(t) = \log(t)$. The Hurst parameter H in (1) is in the range $1/2 < H < 1$ and it characterizes the process in terms of the degree of self-similarity and long-range dependence. The degree of self-similarity and long-range dependence increases as $H \to 1$. For each $m = 1, 2, 3, \cdots$, let $X^{(m)} = \{X_k^{(m)} : \ k = 1, 2, 3, \cdots\}$ denote a new time series obtained by averaging the original series X over non-overlapping blocks of size m; that is, for each $m = 1, 2, 3, \cdots$, $X_k^{(m)}$ is given by

$$X_k^{(m)} = \frac{X_{km-m+1} + \cdots + X_{km}}{m},$$
$$k = 1, 2, 3, \cdots$$

Note that for each m, the aggregate time series $X^{(m)}$ defines a covariance stationary process.

The notion of *self-similarity* and *asymptotic self-similarity* has been introduced in the literature (see, e.g., [4], [13], [15], [20]). A process X is exactly or asymptotically second-order self-similar if the corresponding aggregate processes $X^{(m)}$ are the same as X, or, become indistinguishable from X at least with respect to their autocorrelation functions. The fractional Gaussian noise (FGN) [13] and the fractional Brownian motion [11] are examples of exactly self-similar processes and the fractional autoregressive integrated moving average process is an example of asymptotically self-similar processes [3].

The simplest models with long-range dependence are self-similar processes, which are characterized by hyperbolically decaying autocorrelation functions. Self-similar and asymptotically self-similar processes are particularly attractive traffic models because the long-range dependence can be characterized by a single parameter H. Self-similar traffic models offer parsimonious descriptions of complex traffic processes, though the complete analysis of these models is an area where further research is needed [7]. Self-similarity manifests itself in a couple of different ways: A spectral density that diverges at the origin [$f(\lambda) \sim 1/\lambda^{\alpha}, \quad 0 < \alpha < 1$], and a non-summable autocorrelation function (indicating long-range dependence) [4]. Another characteristic that spans many time scales is the heavy-tailed nature of the density function describing self-similar processes [5].

In [14], the FGN is introduced as a family of random processes that the interdependence between values of the process at instants of time far distant from each other is small but non-negligible. The term "fractional noise" can be explained from spectral theory. Classically, a white noise is defined as a random process having a spectral density independent of the frequency f. As a result, the time integral and derivative of the white noise, and its repeated integrals or derivatives, all have spectral densities of the form $f^{-2\beta}$, with β an integer. Fractional noise, on the contrary, can be defined as having a spectral density of the form $f^{-2\beta}$, with β a non-integer fraction. This explains the term

"fractional noise," or more precisely "fractional white noise." By repeated integration or differentiation, one can restrict β to any interval of unit length. It can be shown that for the FGN, $\beta = H + 0.5$ [14].

In the present paper, we adopt the notation in [2], where the power spectrum of the FGN is given by

$$f(\lambda; H) = A(\lambda; H)\left[|\lambda|^{-2H-1} + B(\lambda; H)\right] \tag{3}$$

for $0 < H < 1$ and $-\pi \leq \lambda \leq \pi$, where

$$A(\lambda; H) = 2\sin(\pi H)\,\Gamma(2H+1)(1-\cos\lambda)$$

and

$$B(\lambda; H) = \sum_{k=1}^{\infty}\Big[(2\pi k + \lambda)^{-2H-1} + (2\pi k - \lambda)^{-2H-1}\Big] \tag{4}$$

In the above expression for $A(\lambda; H)$, $\Gamma(\cdot)$ indicates the Gamma function. Simple calculations show that the FGN is exactly second-order self-similar with self-similarity parameter H, as long as $0.5 < H < 1$ [7], [13].

3 Computation of the FGN Power Spectrum Using a Linear Approximation

There are several existing methods reported in the literature for synthesizing sample paths of self-similar processes including the FGN. The existing methods include *autoregressive processes* [13], [20], *fractional autoregressive integrated moving average process* [3], [4], [9], *fast Fourier transform* (FFT) [15], *alternating renewal process* [21], *queueing processes* [13], [16], [20], *random midpoint displacement* [12], *spatial renewal process* [18], and *wavelet transformation* [8].

An approach is proposed in [15] for generating the FGN using FFT. Suppose that the power spectrum $f(\lambda; H)$ of the FGN process in (3) is known. Then, we can construct a sequence of complex numbers corresponding to this power spectrum, i.e., construct a frequency-domain sample path. An inverse Fourier transform can then be used to obtain the time-domain counterpart sequence. The main difficulty with this approach lies in accurately computing $f(\lambda; H)$ in a timely manner. Another problem is to find a frequency-domain sample path that truly corresponds to the FGN power spectrum. In [15], an approximation is performed by using the midpoint value of two integrals and by retaining some terms of the infinite summation that is used to calculate the FGN power spectrum.

As we can see from equation (4), each term in equation (4) is a function of the index k of the summation and the frequency λ. The calculation of the power spectrum of the FGN implies the computation of an infinite summation for every frequency λ. In order to find a fast and accurate way to approximate equation (4), we need to analyze the terms in the summation. Usually, λ takes values in the range from 0 to π due to symmetry of the spectrum. The number of frequency samples depends on the length of the trace required. The best way to evaluate the summation consists of exploiting the speed improvement of vector operations offered by the programming languages. In this case, λ is a vector which contains the different discrete frequency values.

Consider the first term in (4),

$$(2\pi k + \lambda)^{-2H-1}.$$

This term decreases as k goes to infinity. Actually, for big values of k this term remains

approximately constant for all the values of λ in the range $0 \leq \lambda \leq \pi$. Table 1 shows the relative errors between $(2\pi k + \lambda)^{-2H-1}$ and its linear approximation in the form of $p\lambda + q$ for several values of k and λ, where p and q are determined by minimizing a mean-squared error. We can see from the table that the error is much larger for $k = 1$ than that for $k > 1$. As the value of k increases, $(2\pi k + \lambda)^{-2H-1}$ is approximately a linear function of λ. For $k \geq 3$ the deviation is very small. Table 2 shows the relative errors between the second term in (4), $(2\pi k - \lambda)^{-2H-1}$, and its linear approximation. We can draw conclusions similar to the above from Table 2.

Table 1: Deviation of $(2\pi k + \lambda)^{-2H-1}$ from linearity ($H = 0.75$).

λ	$k = 1$	$k = 2$	$k = 3$	$k = 4$
0.314	6.03%	2.14%	1.09%	0.65%
0.628	1.64%	0.65%	0.34%	0.21%
0.943	-1.87%	-0.51%	-0.23%	-0.13%
1.257	-4.36%	-1.31%	-0.62%	-0.36%
1.571	-5.66%	-1.72%	-0.82%	-0.48%
1.885	-5.61%	-1.72%	-0.82%	-0.48%
2.199	-4.06%	-1.26%	-0.61%	-0.36%
2.513	-0.81%	-0.33%	-0.17%	-0.11%
2.827	4.31%	1.10%	0.49%	0.28%
π	11.48%	3.06%	1.39%	0.79%

As shown in Tables 1 and 2, each term in (4) ($k \geq 3$) can be approximated using a linear function of λ. Because the addition of linear functions is also a linear function, it is possible to use a linear approximation to estimate the infinite sum in equation (4). The illustration in Tables 1 and 2 suggests that a better approximation using a linear function for B in (4) should not include the terms for $k = 1$ and $k = 2$. Figure 1 shows $B_{2:\infty}$ for different values of H. The index $2{:}\ \infty$ means that the summation was performed starting with $k = 2$ and ending at a k that the remaining error in the computation of $B_{2:\infty}$ is less than 0.0001%. In this case, $B_{2:\infty}$ is almost a linear function of λ. For $0.7 < H < 1$ the linear approximation does a better job than for $0.5 < H < 0.7$. On the other hand, $B_{3:\infty}$ can be perfectly approximated, for any H, using a linear function of λ, as illustrated in Figure 2. In the sequel, we will develop a method for using a linear approximation to evaluate equation (4).

Table 2: Deviation of $(2\pi k - \lambda)^{-2H-1}$ from linearity ($H = 0.75$).

λ	$k = 1$	$k = 2$	$k = 3$	$k = 4$
0.314	49.21%	5.68%	2.07%	1.06%
0.628	19.70%	2.08%	0.37%	0.37%
0.943	-0.75%	-0.52%	-0.24%	-0.14%
1.257	-13.35%	-2.19%	-0.88%	-0.47%
1.571	-19.24%	-3.01%	-1.19%	-0.63%
1.885	-19.51%	-3.03%	-1.19%	-0.63%
2.199	-15.23%	-2.32%	-0.91%	-0.48%
2.513	-7.37%	-0.94%	-0.35%	-0.18%
2.827	3.13%	1.03%	0.47%	0.27%
π	15.4%	3.54%	1.53%	0.85%

Using (3) to get the power spectrum involves the infinite summation as in (4). We will now write (4) as

$$B(\lambda; H) = \sum_{k=1}^{2} \Big[(2\pi k + \lambda)^{-2H-1} + (2\pi k - \lambda)^{-2H-1} \Big] + B_{3:\infty}$$

where

$$B_{3:\infty} = \sum_{k=3}^{\infty} \Big[(2\pi k + \lambda)^{-2H-1} + (2\pi k - \lambda)^{-2H-1} \Big]$$

In order to approximate $B_{3:\infty}$ using a linear function of λ, we define $D(\lambda; H)$ as

$$D(\lambda; H) = p\lambda + q \tag{5}$$

To show the dependence of p and q on H, we should express them as $p(H)$ and $q(H)$, respectively. Nevertheless, to avoid using cumbersome notation, we will use p and q without forgetting their dependence on H. We will determine p and q through optimization in the sense of a mean-squared error. We define the mean-squared error ε as

$$\varepsilon = \int_0^{\pi} \left[B_{3:\infty}(\lambda; H) - D(\lambda; H)\right]^2 d\lambda \tag{6}$$

The integral goes from 0 to π because we need to minimize the error only in that interval.

The minimum value of ε is obtained when $\partial\varepsilon/\partial p = 0$ and $\partial\varepsilon/\partial q = 0$. Hence, using Leibnitz's rule [17] to derive under the integral sign, we get

$$\frac{\partial \varepsilon}{\partial q} = \int_0^{\pi} 2\Bigg[\sum_{k=3}^{\infty} \Big((2\pi k + \lambda)^{-2H-1} + (2\pi k - \lambda)^{-2H-1} \Big) - (p\lambda + q) \Bigg] (-1) d\lambda$$

The summation inside the integral defines a function which is continuous on the closed and bounded interval [0, π]. We can interchange the order of the summation and the integral [6] to obtain

$$\frac{\partial \varepsilon}{\partial q} = -2 \sum_{k=3}^{\infty} \Bigg[\int_0^{\pi} (2\pi k + \lambda)^{-2H-1} d\lambda + \int_0^{\pi} (2\pi k - \lambda)^{-2H-1} d\lambda \Bigg] + 2 \int_0^{\pi} (p\lambda + q) d\lambda$$

By letting $\partial\varepsilon/\partial q = 0$, we get

$$\sum_{k=3}^{\infty} \Bigg[\int_0^{\pi} (2\pi k + \lambda)^{-2H-1} d\lambda + \int_0^{\pi} (2\pi k - \lambda)^{-2H-1} d\lambda \Bigg] = \int_0^{\pi} (p\lambda + q) d\lambda$$

which implies that

$$\frac{\pi^2}{2} p + \pi q = F(H) \tag{7}$$

where $F(H)$ is given by

$$F(H) = \sum_{k=3}^{\infty} \left[\frac{(2\pi k - \pi)^{-2H} - (2\pi k + \pi)^{-2H}}{2H} \right] \tag{8}$$

We proceed in the same manner to calculate $\partial\varepsilon/\partial p$. By letting $\partial\varepsilon/\partial p = 0$, we get

$$\sum_{k=3}^{\infty} \Bigg[\int_0^{\pi} \lambda (2\pi k + \lambda)^{-2H-1} d\lambda + \int_0^{\pi} \lambda (2\pi k - \lambda)^{-2H-1} d\lambda \Bigg] = \int_0^{\pi} (p\lambda^2 + q\lambda) d\lambda$$

After some manipulations, we get

$$\frac{\pi^3}{3} p + \frac{\pi^2}{2} q = G(H) \tag{9}$$

where $G(H)$ is given by

$$G(H) = \sum_{k=3}^{\infty} \Bigg[\Big((2\pi k)(2\pi k + \pi)^{-2H} + (2\pi k)(2\pi k - \pi)^{-2H} - 2(2\pi k)^{-2H+1} \Big) / 2H - \Big((2\pi k + \pi)^{-2H+1} + (2\pi k - \pi)^{-2H+1} - 2(2\pi k)^{-2H+1} \Big) / (2H - 1) \Bigg] \tag{10}$$

We note that (10) is for the case when $H > 0.5$. For comparison studies in Section 4, we also derived $G(H)$ for $H = 0.5$ as

$$G(H) = \sum_{k=3}^{\infty} \left[\frac{2\pi k}{2\pi k + \pi} + \frac{2\pi k}{2\pi k - \pi} \right.$$

$$+\ln(2\pi k+\pi)+\ln(2\pi k-\pi)$$

$$-2-2\ln(2\pi k)\Big] \qquad (11)$$

Solving for p and q in equations (7) and (9), we get

$$p=-\frac{6}{\pi^2}F(H)+\frac{12}{\pi^3}G(H) \qquad (12)$$

$$q=\frac{4}{\pi}F(H)-\frac{6}{\pi^2}G(H) \qquad (13)$$

Equations (12) and (13) are the core of our approximation. Figure 3 shows plots of the functions p_n versus the number of terms used in the summation (n) for several values of H. The slope of the linear approximation, p, converges very fast, and a value of $n=20$ seems to be sufficient. The illustration for function q shows that q does not converge that fast, and a minimum value of $n=200$ is required. We also performed experiments for q_n using $n=400$, which resulted in no significant improvements.

We note that in the computation of (3) and (4) for the power spectrum of the FGN, infinite summations are required for every λ. In the present approach, using the linear approximation for $B_{3:\infty}$ as in (5), we do not need to perform infinite summations for any λ. Instead, we only need to perform two infinite summations as in (8) and (10)/(11) for the computation of parameters (p,q) in the linear approximation. This is a significant computational reduction.

The approximation in [15], as described earlier, utilizes the midpoint value of two integrals and retains some terms of the infinite summation for the calculation of the FGN power spectrum. In comparison, the present approach utilizes linear approximation for the infinite summation of the FGN power spectrum. It is emphasized that in both the present approach and that of [15], an inverse Fourier transform (i.e., FFT) is performed to obtain the time-domain sequence. We will perform a comparison study of the present approach with that of [15] in the next section.

4 Simulation Studies

In this section, we will compare the computation time in generating self-similar network traffic based on the FGN, the fractional Brownian motion, and the fractional autoregressive integrated moving average process, and compare the accuracy of the traffic generated with given Hurst parameters.

To compare first the approximation error of the present approach with [15], we calculated the relative error of our approach for the computation of $B(\lambda;H)$. The relative error is defined as

$$E_r=\frac{B(\lambda;H)-[B_{1:2}(\lambda;H)+p\lambda+q]}{B(\lambda;H)}$$

$$=\frac{B_{3:\infty}(\lambda;H)-(p\lambda+q)}{B(\lambda;H)} \qquad (14)$$

In the calculation of $B(\lambda;H)$ in (14), we used 10,000 terms in the summation which is the same number of terms as used in [15]. In all our simulation studies, a Sun SPARC 20 workstation and Matlab were used to run our algorithms. As we can see from Figure 4, the relative error does not exhibit any bias and it is almost independent of H with a maximum value of 0.15%. On the other hand, it has been reported in [15] that Paxson's approximation exhibits an error up to 0.5% with a positive bias.

There are three different methods to estimate the Hurst parameter of self-similar processes [4]: (1) time-domain analysis based on the R/S statistics; (2) frequency-domain

analysis based on the periodogram; and (3) analysis of variances of the aggregate processes, i.e., the variance-time plot method. We conduct in this paper a comparison study of the three methods for estimating the Hurst parameter of sequences generated using our approach. The results are shown in Table 3. For each value of $H = 0.50, 0.55, \cdots, 0.95$, we synthesized 100 sample series of length 32,768 using the present approach. Then, we used the above three methods to estimate the value of H. Table 3 shows the mean (the first row) and the standard deviation (the second row) for each value of H estimated using each method. As we can see from Table 3, the variance-time plot method, with a maximum deviation 9%, has the smallest deviation in most cases, while the periodogram-based method shows the smallest bias among the three methods. For aggregation of traffic generated using the present approach at levels of $m = 1, 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024$, we estimated the Hurst parameters using the three methods mentioned above. The estimates were extremely stable and practically constant over the range of aggregation levels $1 \leq m \leq 1024$. The standard deviation was below 0.005 for both the variance-time plot method and the periodogram-based method. For the R/S statistics method, the standard deviation was a little higher than the other two methods, and values around 0.07 were observed. Because the range includes small values of m, the synthetic series can be regarded as exactly self-similar [4], [13]. We recall that the FGN is exactly second-order self-similar.

Our next simulation compares the computation time of the present approach and that of [15] in approximating the infinite summation $B(\lambda; H)$ for the computation of the FGN power spectrum. The results are shown in Table 4. We note that in Table 4, we only show the time for the computation of $B(\lambda; H)$, and it does not include the computation time for the inverse FFT, for example. We can see from Table 4 that Paxson's approximation takes about twice as much time as ours in all cases considered. By carefully comparing Paxson's method with ours, we can see that Paxson's method requires more than double vector operations than our method and hence, it requires more computation time and more memory space than our method.

Table 3: Comparison of the three methods for estimating the Hurst parameter.

Target H	Variance-Time	R/S Statistics	Periodogram
0.50	0.500	0.553	0.506
	0.006	0.017	0.004
0.55	0.550	0.595	0.551
	0.006	0.018	0.020
0.60	0.598	0.633	0.604
	0.006	0.021	0.009
0.65	0.650	0.675	0.656
	0.005	0.019	0.009
0.70	0.697	0.712	0.707
	0.006	0.021	0.008
0.75	0.747	0.751	0.759
	0.007	0.023	0.008
0.80	0.793	0.783	0.807
	0.006	0.021	0.008
0.85	0.839	0.819	0.858
	0.007	0.023	0.009
0.90	0.877	0.847	0.908
	0.009	0.019	0.010
0.95	0.912	0.877	0.959
	0.009	0.020	0.007

To compare the total time needed for generating self-similar sequences, we conducted the following simulation study. Table 5 shows the running time of several algorithms for generating self-similar sequences. The second column of Table 5 shows the running time

of our approach based on the linear approximation. This table also shows the running time using Paxson's approximation [15]. In our approach and the approach of [15], an inverse Fourier transform is performed to generate the time sequence for self-similar traffic. Again, because Paxson's method requires more than double vector operations than our method, it requires more computation time and more memory space than our method. The rest of Table 5 shows the running time of the random midpoint displacement method (RMD) for generating the fractional Brownian motion [12], the spatial renewal process (SRP) method for generating the FGN [18], and the Hosking's method for generating the fractional autoregressive integrated moving average process [9]. Each method has advantages and disadvantages. Hosking's method is the slowest one, while our method is the fastest of all the methods in Table 5.

Table 4: Running time in seconds for approximating $B(\lambda; H)$.

Length	Our approach	Paxson's
8,192	0.18	0.37
16,384	0.37	0.80
32,768	0.84	2.06
65,536	1.64	3.38
131,072	3.27	6.84
262,144	6.49	14.97
524,288	13.00	27.12
1,048,576	86.99	168.69
2,097,152	379.36	724.50

The method based on the alternating renewal process [21] requires the multiplexing of an infinite number of independent process. This is one of its disadvantages, and a balance between speed and quality of the generated sequence must be determined. The SRP method has the same disadvantage. It requires a minimum aggregation level of ten. The running times shown in Table 5, for the SRP model, were calculated using an aggregation of ten independent processes; the probability distribution was approximated using a piecewise function of five linear functions. The quality of the generated traffic using the SRP is not satisfactory, and the value of the Hurst parameter does not remain constant for several levels of aggregation, which agrees with the findings in [18]. The RMD method is one of the fastest, as shown in Table 5. Nevertheless, it fails to provide stationary increments when $H \neq 0.5$ which implies that the resulting process is only an approximation of FBM [8]. Furthermore, the wavelet interpretation of the RMD method indicates that the system is not orthonormal [8]. The FFT method, using Paxson's approximation and the present approach based on a linear approximation, generates high quality self-similar sequences, where the Hurst parameters of the samples agree with the target values. We observed that the Hurst parameter remains constant for several levels of aggregation in the variance-time plot, in the R/S statistics, and in the periodogram.

Finally, we use Figure 5 to show a few synthetic self-similar series generated using the present method. From this figure, we can clearly see that strongly correlated data are generated for large values of H. For large H, the series tends to exhibit low frequency cycles, which is a typical behavior of self-similar/LRD traces. It is also observed in the figure that for values of H close to 0.5 (short-range dependence) the sequence resembles white noise, as expected.

It is worthwhile to say that using the linear approximation method presented in this paper to calculate the FGN power spectrum allows us to quickly generate high quality self-similar sequences.

Table 5: Running time in seconds for generating self-similar sequences.

Length	Our approach	Paxson
65,536	5	7
131,072	10	12
262,144	22	24
524,288	45	48
1,048,576	94	260
2,097,152	259	659

Table 5 (Continued)

Length	RMD	SRP	Hosking
65,536	32	1,306	5,826
131,072	66	4,575	24,152
262,144	130	16,510	104,214
524,288	262	62,213	440,361
1,048,576	530	245,140	too long
2,097,152	1,100	949, 172	too long

5 Conclusions

Long-range dependence is often encountered in practice, not only in hydrology and geophysics, but in all fields of statistical applications including traffic engineering. If not taken into account, it can completely invalidate statistical inference [1], [4]. For many situations, new statistical methods as well as properties of classical techniques are sufficiently known nowadays to be used in practice. From a practical point of view, fast and accurate approximations for generating synthetic self-similar traffic are needed.

We developed in the present paper a new approach for approximating the power spectrum of the fractional Gaussian noise (FGN). To deal with the infinite summation of the FGN spectral density function, we computed the first two terms exactly, and approximated the rest of the summation using a linear function. We illustrated the validity of our approach. We also provided simulation results to compare the present approach with existing ones in terms of the computation speed and the accuracy in generating self-similar traffic with given Hurst parameters. We concluded that our approach is the fastest in computation speed among all the methods in our simulation and our approach guarantees to generate self-similar traffic with high degree of accuracy in terms of burstiness measures.

References

[1] A. Adas and A. Mukherjee, "On resource management and QoS guarantees for long range dependent traffic," *Proceedings of the IEEE Conference on Computer Communications*, Boston, MA, Apr. 1995, pp. 779–787.

[2] J. Beran, *Estimation, Testing and Prediction for Self-Similar and Related Processes*, Ph.D. Dissertation, Swiss Federal Institute of Technology (ETH), Zürich, Switzerland, 1986.

[3] J. Beran, "Statistical methods for data with long-range dependence," *Statistical Science*, vol. 7, pp. 404–416, 1992.

[4] J. Beran, R. Sherman, M. S. Taqqu, and W. Willinger, "Long-range dependence in variable-bit-rate video traffic," *IEEE Transactions on Communications*, vol. 43, pp. 1566–1579, Feb./Mar./Apr. 1995.

[5] M. E. Crovella and A. Bestavros, "Self-similarity in World Wide Web traffic: Evidence and possible causes," *IEEE/ACM Transactions on Networking*, vol. 5, pp. 835–846, Dec. 1997.

[6] C. H. Edwards and D. E. Penney, *Calculus with Analytic Geometry*, Upper Saddle River, NJ: Prentice-Hall, 1998.

[7] A. Erramilli, O. Narayan, and W. Willinger, "Experimental queueing

analysis with long-range dependent packet traffic," *IEEE/ACM Transactions on Networking*, vol. 4, pp. 209–223, Apr. 1996.

[8] P. Flandrin, "Wavelet analysis and synthesis of fractional Brownian motion," *IEEE Transactions on Information Theory*, vol. 38, pp. 910–917, Mar. 1992.

[9] M. W. Garret and W. Willinger, "Analysis, modeling and generation of self-similar VBR video traffic," *Computer Communications Review*, vol. 24, pp. 269–280, Oct. 1994.

[10] D. P. Heyman and T. V. Lakshman, "What are the implications of long-range dependence for VBR-video traffic engineering?" *IEEE/ACM Transactions on Networking*, vol. 4, pp. 301–317, June 1996.

[11] K. R. Krishnan, "A new class of performance results for a fractional Brownian traffic model," *Queuing Systems*, vol. 22, pp. 277–285, 1996.

[12] W.-C. Lau, A. Erramilli, J. L. Wang, and W. Willinger, "Self-similar traffic generation: The random midpoint displacement algorithm and its properties," *Proceedings of the IEEE International Conference on Communications*, Seattle, WA, June 1995, vol. 1, pp. 466–472.

[13] 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, pp. 1–15, Feb. 1994.

[14] B. Mandelbrot and J. R. Wallis, "Computer experiments with fractional Gaussian noises–Part 1: Averages and variances," *Water Resources Research*, vol. 5, pp. 228–267, Feb. 1969.

[15] V. Paxson, "Fast, approximate synthesis of fractional Gaussian noise for generating self-similar network traffic," *Computer Communications Review*, vol. 27, pp. 5–18, Oct. 1997.

[16] V. Paxson and S. Floyd, "Wide area traffic: The failure of Poisson modeling," *IEEE/ACM Transactions on Networking*, vol. 3, pp. 226–244, June 1995.

[17] M. R. Spiegel and J. Liu, *Schaum's Mathematical Handbook of Formulas and Tables*, Second Edition, New York, NY: McGraw-Hill, 1998.

[18] T. Taralp, M. Devetsikiotis, and I. Lambadaris, "Efficient fractional Gaussian noise generation using the spatial renewal process," *Proceedings of the IEEE International Conference on Communications*, Atlanta, GA, June 1998, vol. 3, pp. 1456–1460.

[19] W. Willinger, "Traffic modeling for high-speed networks: Theory versus practice," in *Stochastic Networks* (vol. 71, pp. 395-409), Edited by F. P. Kelly and R. J. Williams, New York, NY: Springer-Verlag, 1995.

[20] W. Willinger, M. S. Taqqu, W. E. Leland, and D. V. Wilson, "Self-similarity in high-speed packet traffic: Analysis and modeling of Ethernet traffic measurements," *Statistical Science*, vol. 10, pp. 67–85, 1995.

[21] W. Willinger, M. S. Taqqu, R. Sherman, and D. V. Wilson, "Self-similarity through high-variability: Statistical analysis of Ethernet LAN traffic at the source level," *IEEE/ACM Transactions on Networking*, vol. 5, pp. 71–86, Feb. 1997.

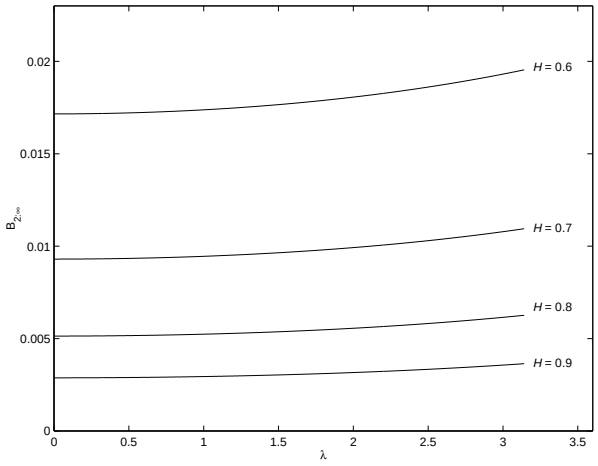

Figure 1: $B_{2:\infty}$ for different values of H and λ

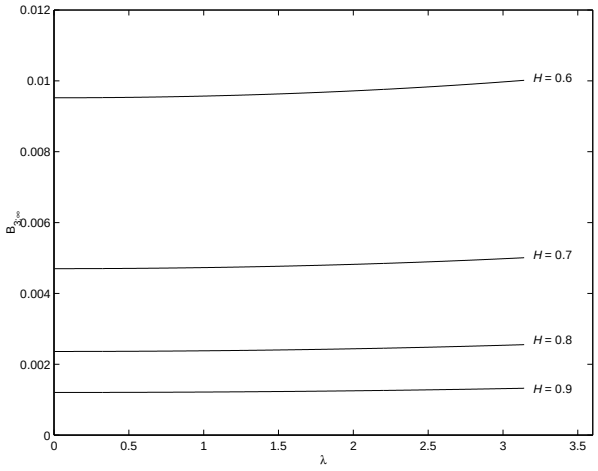

Figure 2: $B_{3:\infty}$ for different values of H and λ

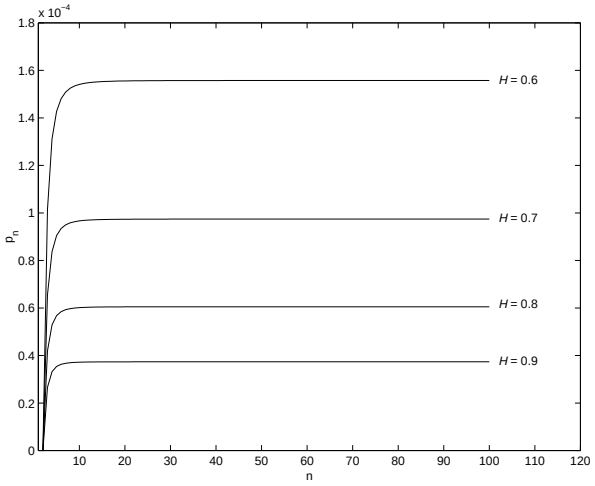

Figure 3: p as a function of n.

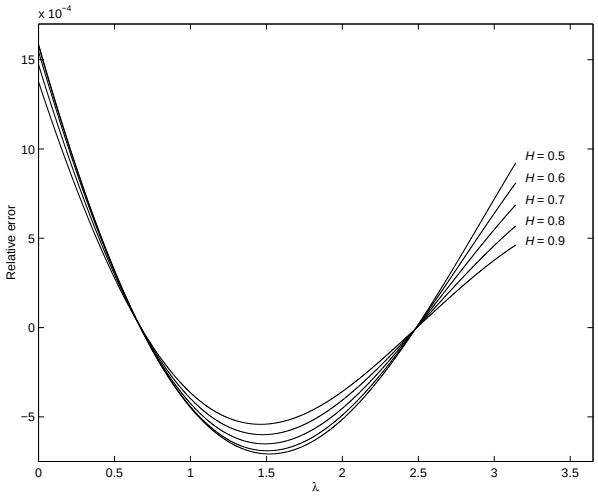

Figure 4: Relative error for several values of H.

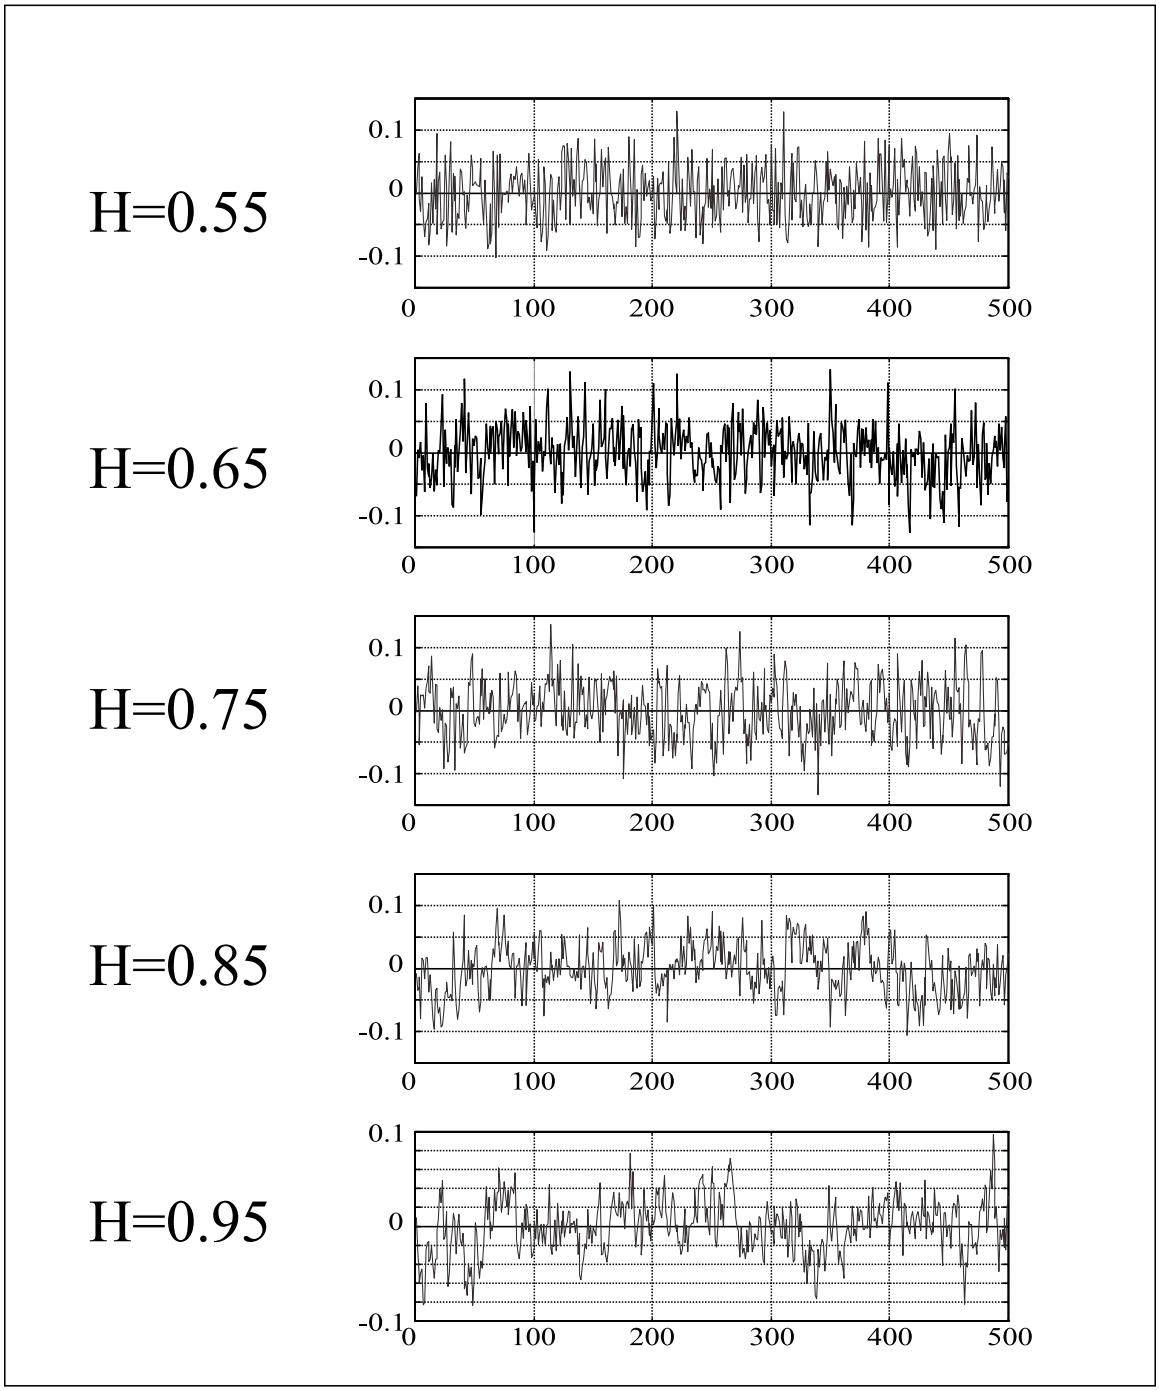

Figure 5: Traces generated using the present approach for several values of H.