

On spectral simulation of fractional Brownian motion

A. B. Dieker* and M. Mandjes

CWI

P.O. Box 94079

1090 GB Amsterdam, the Netherlands

and

University of Twente

Faculty of Mathematical Sciences

P.O. Box 217

7500 AE Enschede, the Netherlands

{Ton.Dieker, Michel.Mandjes}@cwi.nl

Abstract

This paper focuses on simulating fractional Brownian motion (fBm). Despite the availability of several exact simulation methods, attention has been paid to approximate simulation (i.e., the output is approximately fBm), particularly because of possible time savings. In this paper, we study the class of approximate methods that are based on the spectral properties of fBm's stationary incremental process, usually called fractional Gaussian noise (fGn). The main contribution is a proof of asymptotical exactness (in a sense that is made precise) of these spectral methods. Moreover, we establish the connection between the spectral simulation approach and a widely used method, originally proposed by Paxson, that lacked a formal mathematical justification. The insights enable us to evaluate the Paxson method in more detail. It is also shown that spectral simulation is related to the fastest known exact method.

1 Introduction

Fractional Brownian motion (fBm) is a widely used Gaussian process with a variety of applications, e.g. in communications engineering, (geo)physics, finance, bioengineering and fractal imaging. Therefore, simulation of fBm has drawn a lot of attention. Although it is possible to generate a discrete-time realization of fBm ('exact' simulation), many 'approximate' methods have been proposed as alternatives to exact simulation. Each of these methods has its own advantages and drawbacks.

Among the approximate simulation methods, *spectral techniques* have become quite popular in the recent past. Spectral methods can be used for simulating stationary processes. In the Fourier filtering method by Saupe [14, pp. 93–94], the spectral approach is used for approximating fBm. By construction, the resulting sample is a good approximation for fBm in the spectral domain. However, some problems arise the time domain; fBm is a non-stationary process, while the approximating Saupe sample is stationary. This problem is circumvented in the spectral method by Yin [17] by simulating the *stationary* increments of fBm, known as fractional Gaussian noise (fGn). Yin's approach is closely related to the family of spectral simulation methods that is analyzed in the present paper — see also [1, 9].

*Part of this work was done while A. B. Dieker was with the Vrije Universiteit Amsterdam. His research was supported by the Netherlands Organization for Scientific Research (NWO) under grant 631.000.002.

The connections between spectral and related methods are not so transparent, and it is in some cases not even clear why the produced samples should be a *good* approximation for realizations of fGn or fBm. Moreover, it is very difficult to develop a fair criterion for evaluation of the approximate methods; one must trade off accuracy and efficiency (speed). Most accuracy analyses include the statistical estimation or testing of the samples [11, 12]. As we will see, these tests may be misleading.

In this paper, we show that the spectrally simulated samples are (in some sense) asymptotically exact. Moreover, we establish the links between the spectral approach and the method proposed by Paxson [12]. This enables us to conclude that Paxson samples are also asymptotically exact. In addition, it makes it possible to study the mean square error at each sample point and the covariance structure of the approximating sample.

The paper is organized as follows. Section 2 consists of the preliminaries: the introduction of fBm and fGn, and a brief review of exact generation of fGn. Since fGn is the incremental process of fBm, a discretized realization of fBm is then obtained from this sample by taking cumulative sums. Some background on spectral densities is also provided in Section 2. The main results of this paper are presented in Section 3. The intuitively plausible fact is proved that spectrally simulated samples are (in a sense that will be made precise) asymptotically exact, and the connections between the different spectral methods are established. The error analysis and the numerical comparison of the approximate covariances are the topic of Section 4.

2 Preliminaries

2.1 Definitions

In the pioneering work of Mandelbrot and van Ness [10], fBm was originally defined as a stochastic integral with respect to ordinary Brownian motion. The family of fractional Brownian motions is indexed by a single parameter $H \in (0, 1)$, which is called the *self-similarity* or *Hurst parameter*. Standardized fBm with Hurst parameter H is a centered continuous-time Gaussian process $B_H(\cdot)$ with covariance function

$$\rho(s, t) \equiv \mathbb{E}B_H(s)B_H(t) = \frac{1}{2} \{t^{2H} + s^{2H} - |t - s|^{2H}\} \quad (1)$$

for $s, t \geq 0$. Note that B_H reduces to an ordinary Brownian motion for $H = 1/2$.

The incremental process of fractional Brownian motion is a stationary discrete-time process and is called *fractional Gaussian noise* (fGn). We define the fractional Gaussian noise $X = \{X_k : k = 0, 1, \dots\}$ by

$$X_k := B_H(k + 1) - B_H(k).$$

It is clear that any X_k has a standard normal distribution; the complication, however, is the covariance between the X_k . To be precise, straightforward computations show that the autocovariance function $\gamma(\cdot)$ of X is given by

$$\gamma(k) \equiv \mathbb{E}X_{n+k}X_n = \frac{1}{2} [|k - 1|^{2H} - 2|k|^{2H} + |k + 1|^{2H}] \quad (2)$$

for $k \in \mathbb{Z}$. By writing down the Taylor expansion of the function $h(x) = (1 - x)^{2H} - 2 + (1 + x)^{2H}$ at the origin and noting that $\gamma(k) = \frac{1}{2}k^{2H}h(1/k)$ for $k \geq 1$, it is easily seen from (2) that

$$\gamma(k) \sim H(2H - 1)|k|^{2H-2} \quad (3)$$

as $|k| \rightarrow \infty$. Therefore, the autocovariance function is non-summable for $H > 1/2$. This phenomenon is called *long-range dependence*, indicating (relatively) slow decay of the covariance function. The present paper covers both the long-range dependent case, i.e. $H > 1/2$, and the short-range dependent case $H \leq 1/2$. However, we often focus on the long-range dependent case, since it is the worst case for the spectral approach (as will become clear in Section 3).

2.2 Exact simulation of fractional Brownian motion

There exist algorithms for simulating general stationary Gaussian processes with a given autocovariance function. Evidently, these algorithms can be used for generating fGn. An fBm sample is then found by computing the cumulative sums of this fGn sample. The so-called self-similarity property (i.e., the fact that $B_H(at)$ has the same finite-dimensional distributions as $a^H B_H(t)$) can be used to obtain a sample on an arbitrary equispaced grid.

The method proposed by Hosking [7] simulates the sample recursively, conditioning on all so far generated points. The required computation time is of order N^2 when N sample points are needed. The Cholesky method [1] uses the Cholesky decomposition of the covariance matrix, which results in an order N^3 algorithm. As pointed out in [5], both methods implicitly compute the same matrix, but the Hosking method achieves that more efficiently. However, the Cholesky method can also be applied to the covariance matrix of fBm, in which case the methods work differently.

In the specific case of fGn, another exact method is the method first mentioned by Davies and Harte [4]. It is of order $N \log N$, i.e., even faster than the Hosking method. The method was later simultaneously generalized by Dietrich and Newsam [6] and Wood and Chan [16]. The algorithm is based on the fact that the covariance matrix of a stationary discrete-time Gaussian processes can be embedded in a so-called circulant matrix. This latter matrix should be positive definite for the algorithm to work, which is indeed the case for fGn. The constructed circulant matrix can be diagonalized explicitly, and the computations are done efficiently with the so-called Fast Fourier Transform (FFT) algorithm. To be able to establish the connections between this method and the spectral simulation approach of Section 3, we describe the algorithm in more detail.

With the Davies and Harte method, an fGn sample of size N can be constructed as follows:

- Define $\alpha_k := \gamma(k)$ for $k = 0, \dots, N-1$, $\alpha_N := 0$ and $\alpha_k := \gamma(2N-k)$ for $k = N+1, \dots, 2N-1$. Compute the (discrete) Fourier transform $(\lambda_k)_{k=0}^{2N-1}$ of $(\alpha_k)_{k=0}^{2N-1}$ with the FFT algorithm. Recall that the Fourier transform of $(\alpha_k)_{k=0}^{j-1}$ is given by

$$\lambda_n = \sum_{k=0}^{j-1} \alpha_k \exp\left(2\pi i \frac{nk}{j}\right) \quad (4)$$

for $n = 0, \dots, j-1$. When j is a power of two, the number of calculations required by the FFT algorithm is of order $j \log_2(j)$; a considerable gain in speed compared to the straightforward calculation of order j^2 . The λ_j are real by construction and they are non-negative in the fGn case when α_N is set to $\gamma(N)$, see [3]. However, we set $\alpha_N = 0$ for reasons that become clear in Section 3. This has no influence on the resulting fGn sample as long as the λ_j are non-negative (in practice, this is satisfied for any reasonable sample size N).

- The output fGn sample is Z_0, \dots, Z_{N-1} , where the sequence $(Z_k)_{k=0}^{2N-1}$ is the Fourier transform of

$$w_k := \begin{cases} \sqrt{\frac{\lambda_k}{2N}} U_k^{(0)} & k = 0; \\ \sqrt{\frac{\lambda_k}{4N}} (U_k^{(0)} + iU_k^{(1)}) & k = 1, \dots, N-1; \\ \sqrt{\frac{\lambda_k}{2N}} U_k^{(0)} & k = N; \\ \sqrt{\frac{\lambda_k}{4N}} (U_{2N-k}^{(0)} - iU_{2N-k}^{(1)}) & k = N+1, \dots, 2N-1, \end{cases} \quad (5)$$

with $U_k^{(i)}$ i.i.d. standard normal random variables for $k = 0, \dots, 2N-1$, and mutually independent vectors $U^{(0)}$ and $U^{(1)}$.

Note that Z_N, \dots, Z_{2N-1} is also an fGn sample, but this sample is not independent of Z_0, \dots, Z_{N-1} , nor is Z_0, \dots, Z_{2N-1} an fGn sample. The reader may easily check that Z is real by construction.

In addition to these exact methods, several approximation algorithms have been proposed. Some of them can be found in [5], including spectral simulation. The Conditionalized Random Midpoint Displacement method [11] is notable for its accuracy, but we remark that the latter method does not produce stationary samples to approximate fGn.

2.3 Spectral densities

Instead of analyzing a stochastic process in the time domain, one could consider the so-called *spectral* or *frequency* domain. It turns out that the so-called spectral density characterizes all frequency information of stationary processes (see e.g. [13]). This spectral density is computed as follows for frequencies $-\pi \leq \lambda \leq \pi$:

$$f(\lambda) := \sum_{j=-\infty}^{\infty} \gamma(j) \exp(ij\lambda), \quad (6)$$

where $\gamma(\cdot)$ denotes the autocovariance function (e.g. of fGn (2)). The autocovariance function is recovered by applying the inversion formula

$$\gamma(j) = \frac{1}{2\pi} \int_{-\pi}^{\pi} f(\lambda) \exp(-ij\lambda) d\lambda. \quad (7)$$

In this paper, we are particularly interested in the spectral density of fractional Gaussian noise. It can be seen [15] that this density is given by

$$f(\lambda) = 2 \sin(\pi H) \Gamma(2H+1) (1 - \cos \lambda) [|\lambda|^{-2H-1} + B(\lambda, H)], \quad (8)$$

where $\Gamma(\cdot)$ denotes the Gamma function and

$$B(\lambda, H) := \sum_{j=1}^{\infty} \{(2\pi j + \lambda)^{-2H-1} + (2\pi j - \lambda)^{-2H-1}\} \quad (9)$$

for $-\pi \leq \lambda \leq \pi$.

Unfortunately, it is not known how this can be simplified further. However, a quite useful result is the proportionality of this spectral density to $|\lambda|^{1-2H}$ near $\lambda = 0$, checked by noting that $1 - \cos(\lambda) = \frac{1}{2}\lambda^2 + O(\lambda^4)$ as $|\lambda| \rightarrow 0$. Therefore, f has a pole at zero for $H > 1/2$. For

$H < 1/2$, f is non-differentiable at zero. From Equation (6), it follows that a pole at zero in the spectral density corresponds to long-range dependence.

To numerically evaluate the spectral density of fGn, the infinite sum in Equation (9) must be truncated. When the truncation parameter is chosen quite large, the function evaluation becomes computationally more demanding. Paxson [12] suggests and tests a useful approximation to overcome this problem. He shows that

$$\tilde{B}_3(\lambda, H) := \sum_{j=1}^3 \left\{ (a_j^+)^{-2H-1} + (a_j^-)^{-2H-1} \right\} + \frac{(a_3^+)^{-2H} + (a_3^-)^{-2H} + (a_4^+)^{-2H} + (a_4^-)^{-2H}}{8H\pi} \quad (10)$$

approximates $f(\lambda)$ quite well, where $a_j^\pm = 2\pi j \pm \lambda$.

3 Spectral simulation, the Paxson method, and the approximate circulant method

The idea of spectral simulation is to simulate a process in the spectral domain and transform the resulting sample to the time domain. Although it is not possible to obtain exact fGn or fBm samples by taking this approach, we will shortly see that the accuracy increases as the sample size grows.

In the course of the exposition, it becomes clear that spectral simulation is closely related to the Paxson method and the Davies and Harte method. The resulting modification of the Davies and Harte method (which produces approximate samples, but is also faster) will be called the *approximate circulant* method.

3.1 Spectral simulation

The spectral analysis of a stationary discrete-time Gaussian process $X = \{X_n : n = 0, \dots, N-1\}$ shows that it can be represented in terms of the spectral density as

$$X_n = \int_0^\pi \sqrt{\frac{f(\lambda)}{\pi}} \cos(n\lambda) dB_1(\lambda) - \int_0^\pi \sqrt{\frac{f(\lambda)}{\pi}} \sin(n\lambda) dB_2(\lambda), \quad (11)$$

where the equality is to be understood as equality in distribution. The two integrators are mutually independent (ordinary) Brownian motions. This result is called the spectral theorem; using (7), its validity is easily checked by showing that the right-hand side of (11) has covariance function (2). In what follows, we assume that the required sample size N is a power of two.

We briefly interrupt the exposition to make clear how $\int_0^\pi \phi(s) dB(s)$ should be interpreted for some function ϕ and an ordinary Brownian motion B as stochastic integrator. This will be helpful for establishing the result that spectrally simulated samples are (in some sense) asymptotically exact. The integral $\int_0^\pi \phi(s) dB(s)$ has a natural definition when the integrand ϕ is a so-called simple function, which means that there exists a positive integer ℓ and a strictly increasing sequence of real numbers $(t_j)_{j=0}^\ell$ with $t_0 = 0$ and $t_\ell = \pi$, as well as a bounded sequence of real numbers $(\phi_j)_{j=0}^{\ell-1}$ such that $\phi(s)$ can be written as $\phi(s) = \phi_j$ for $s \in (t_j, t_{j+1}]$ (the value at 0 is an arbitrary real number). For such a simple function ϕ , the stochastic integral has the following definition:

$$\int_0^\pi \phi(s) dB(s) := \sum_{j=0}^{\ell-1} \phi_j (B(t_{j+1}) - B(t_j)). \quad (12)$$

It is known that every square integrable function ψ can be written as a limit in L^2 (the space of all Lebesgue square integrable functions on $[0, \pi]$) of a sequence (ψ_m) of simple functions; the subset of simple functions is dense in L^2 . Then there exists a random variable Z such that

$$\lim_{m \rightarrow \infty} \mathbb{E} \left[\left(Z - \int_0^\pi \psi_m(s) dB(s) \right)^2 \right] = 0. \quad (13)$$

This random variable is unique in the sense that another random variable that satisfies (13) is almost sure equal to Z . This makes it possible to take this Z as the definition of the stochastic integral $\int \psi(s) dB(s)$. We refer to the type of convergence (13) as *convergence in \mathcal{L}^2 -norm*. Moreover, it can be shown that Z is independent of the approximating sequence $(\psi_m)_m$. We will need the following implication of this observation: when a given sequence of simple functions $(\psi_m)_m$ converges in L^2 -norm to ψ , the corresponding sequence of stochastic integrals $(\int \psi_m(s) dB(s))_m$ converges in \mathcal{L}^2 -norm to $\int \psi(s) dB(s)$. More on stochastic integration can be found in any textbook on stochastic calculus, e.g. [8].

Spectral simulation is based on approximating Equation (11); the integrand is replaced by a simple function. Define $\xi_n(\lambda) := \sqrt{f(\lambda)/(2\pi)} \cos(n\lambda)$ and fix some integer ℓ . After setting $t_k = \pi k/\ell$ for $k = 0, \dots, \ell - 1$, we define a simple function $\xi_n^{(\ell)}$ on $[0, \pi]$ for $0 \leq n \leq N - 1$ by

$$\xi_n^{(\ell)}(\lambda) := \sqrt{\frac{f(t_1)}{\pi}} \cos(nt_1) \mathbf{1}_{\{0\}}(\lambda) + \sum_{k=0}^{\ell-1} \sqrt{\frac{f(t_{k+1})}{\pi}} \cos(nt_{k+1}) \mathbf{1}_{(t_k, t_{k+1}]}(\lambda). \quad (14)$$

Let $\theta_n^{(\ell)}(\lambda)$ be defined as $\xi_n^{(\ell)}(\lambda)$, but with the cosine terms replaced by sine terms. The first integral in (11) is approximated by $\int_0^\pi \xi_n^{(\ell)}(\lambda) dB_1(\lambda)$, which can be computed using Equation (12). Since a similar approximation can be made for the second integral, we arrive at the following approximation $\hat{X}_n^{(\ell)}$ of X_n :

$$\hat{X}_n^{(\ell)} := \sum_{k=0}^{\ell-1} \sqrt{\frac{f(t_{k+1})}{\ell}} \left[\cos(nt_{k+1}) U_k^{(0)} - \sin(nt_{k+1}) U_k^{(1)} \right], \quad (15)$$

where $U_k^{(i)}$ are again i.i.d. standard normal random variables for $k = 0, \dots, \ell - 1$. The two vectors $U^{(0)}$ and $U^{(1)}$ should also be mutually independent, since B_1 and B_2 are independent as well.

It should be noted that for $H > 1/2$, we approximate the spectral density by functions that have no pole at 0. As already pointed out, this pole is equivalent to long-range dependence. Hence, we approximate a long-range dependent process by a short-range dependent process. Still, we may obtain a sample with a covariance structure that approximates the structure determined by (2) and (3) very well. This issue is explored further in Section 4.

The FFT can be used to calculate (15) efficiently. To this end, we define the sequence $(a_k)_{k=0, \dots, 2\ell-1}$ by

$$a_k := \begin{cases} 0 & k = 0; \\ \frac{1}{2} \left(U_{k-1}^{(0)} + iU_{k-1}^{(1)} \right) \sqrt{f(t_k)/\ell} & k = 1, \dots, \ell - 1; \\ U_{k-1}^{(0)} \sqrt{f(t_k)/\ell} & k = \ell; \\ \frac{1}{2} \left(U_{2\ell-k-1}^{(0)} - iU_{2\ell-k-1}^{(1)} \right) \sqrt{f(t_{2\ell-k})/\ell} & k = \ell + 1, \dots, 2\ell - 1. \end{cases} \quad (16)$$

Using (4), one can check that the Fourier transform of (a_k) is real and equals (15).

Since $\xi_n^{(\ell)}$ approximates ξ_n better for larger ℓ , an interesting question is whether the approximate sample converges to an exact sample as $\ell \rightarrow \infty$. However, different types of convergence of random variables and processes exist; it is at first sight not clear in what sense the approximate sample converges. In the remainder of this subsection, we deal with this convergence issue.

Convergence

We start by deriving the covariance structure of $\hat{X}^{(\ell)}$. From (15) it follows that the covariance between $\hat{X}_m^{(\ell)}$ and $\hat{X}_n^{(\ell)}$ for $n, m = 0, \dots, N-1$ is given by

$$\text{Cov}(\hat{X}_m^{(\ell)}, \hat{X}_n^{(\ell)}) = \sum_{k=0}^{\ell-1} \frac{f(t_{k+1})}{\ell} \cos((m-n)t_{k+1}), \quad (17)$$

which depends only on $m-n$. Hence, the spectral simulation method produces stationary approximate samples. A more detailed numerical analysis of the covariance structure is deferred to Section 4.

It is readily checked with (7) that $\text{Cov}(\hat{X}_m^{(\ell)}, \hat{X}_n^{(\ell)})$ converges to $\gamma(|m-n|)$ as $\ell \rightarrow \infty$. From this fact, it is not difficult to deduce that the finite-dimensional distributions of $\hat{X}^{(\ell)}$ converge in distribution to the corresponding finite-dimensional distributions of X as $\ell \rightarrow \infty$. However, we will not prove this, since we can prove an even stronger convergence result: every sample point $\hat{X}_n^{(\ell)}$ converges in \mathcal{L}^2 -norm to the corresponding exact sample point X_n as $\ell \rightarrow \infty$.

The proof of this fact is based on the definition of the stochastic integral appearing in (11). Because f is integrable, the function ξ_n is certainly square integrable for $0 \leq n \leq N-1$. Recall that the discussion on stochastic integration showed that if the sequence of simple functions $(\xi_n^{(\ell)})_\ell$ satisfies, for fixed $0 \leq n \leq N-1$,

$$\lim_{\ell \rightarrow \infty} \int_0^\pi [\xi_n(\lambda) - \xi_n^{(\ell)}(\lambda)]^2 d\lambda = 0, \quad (18)$$

then $\int_0^\pi \xi_n^{(\ell)}(\lambda) dB_1(\lambda)$ converges in \mathcal{L}^2 -norm to $\int_0^\pi \xi_n(\lambda) dB_1(\lambda)$ as $\ell \rightarrow \infty$. It is indeed true that (18) holds; a similar result holds for the second integrand of (11). A proof of these facts can be found in the appendix. By the independence of B_1 and B_2 , we deduce that every sample point $\hat{X}_n^{(\ell)}$ converges to X_n in \mathcal{L}^2 -norm, i.e., in mean square sense.

Since the ‘error’ $\hat{X}_n^{(\ell)} - X_n$ is a centered Gaussian variable for every n (with a variance that decays to zero in ℓ), we have $\mathbb{E}|\hat{X}_n^{(\ell)} - X_n|^{2p} = C_p (\mathbb{E}|\hat{X}_n^{(\ell)} - X_n|^2)^p$ for some constant $C_p > 0$ depending only on p . Thus, $\hat{X}_n^{(\ell)}$ converges to X_n in \mathcal{L}^p -norm for *every* $p \geq 1$! Because of the normality, this is equivalent to convergence in probability. This convergence of the individual sample points is readily extended to a joint convergence result: the finite-dimensional distributions of $\hat{X}^{(\ell)}$ converge in probability to the corresponding finite-dimensional distributions of X as $\ell \rightarrow \infty$.

An interesting question is at what rate the convergence takes place. By the way the stochastic integral is constructed, this rate is related to the rate of convergence of (18) and its θ_n -counterpart. Since the spectral density is proportional to $|\lambda|^{1-2H}$ near $\lambda = 0$, we observe that by monotone convergence:

$$\sigma_{\hat{X}_n^{(\ell)} - X_n}^2 = \|\xi_n - \xi_n^{(\ell)}\|^2 + \|\theta_n - \theta_n^{(\ell)}\|^2$$

$$\begin{aligned}
&\geq \int_0^{\pi/\ell} \left(\sqrt{\frac{f(\lambda)}{2\pi}} \cos(n\lambda) - \sqrt{\frac{f(\pi/\ell)}{2\pi}} \cos(n\pi/\ell) \right)^2 d\lambda \\
&\sim C_H \int_0^{\pi/\ell} \left[\lambda^{1/2-H} - (\pi/\ell)^{1/2-H} \right]^2 d\lambda \\
&= C_H \pi^{1-2H} \left(\frac{1}{4-4H} - \frac{1}{3/2-H} + \frac{1}{2} \right) \ell^{2H-2},
\end{aligned}$$

where C_H denotes some constant. Thus, the rate of convergence is quite slow.

Since the length of the output sequence of the FFT algorithm applied to (16) must be at least the sample size N , the smallest possible choice for ℓ is $N/2$, although better results are obtained for larger ℓ . For $\ell = N/2$, the spectral simulation approach is closely related to the Paxson method, as will be made clear in the next subsection.

3.2 The Paxson method

Paxson [12] proposes a rather intuitive method for simulating fractional Gaussian noise. By studying the output statistically, he tests if the resulting samples have indeed the desired properties. Unfortunately, the paper lacks a thorough justification of the proposed procedure, and it remains unclear why the obtained sample should be (close to) Gaussian.

However, with the formulas above, it is possible to make the arguments precise. In the Paxson method, the approximate fGn sample is the Fourier transform of

$$b_k := \begin{cases} 0 & k = 0; \\ \sqrt{\frac{R_k f(t_k)}{N}} \exp(i\Phi_k) & k = 1, \dots, N/2; \\ b_{N-k}^* & k = N/2 + 1, \dots, N - 1, \end{cases}$$

where R_k are independent exponentially distributed random variables with mean 1 for $k \geq 1$, and the asterisk denotes the complex conjugate. Besides $\Phi_{N/2}$, which is set to zero, the Φ_k are independent uniformly distributed random variables on $[0, 2\pi]$ for $k \geq 1$, also independent of the R_k . In this case, t_k equals $2\pi k/N$. Note that the obtained sample is real by construction.

Because R_k is exponentially distributed with mean 1, $\sqrt{2R_k} \exp(i\Phi_k)$ has the same distribution as $U_k^{(0)} + iU_k^{(1)}$, where $U_k^{(0)}$ and $U_k^{(1)}$ are independent standard normal random variables in the usual notation. This fact is also used in the well-known Box-Muller algorithm to simulate Gaussian random variables (e.g., [1]).

Let us now compare this with the spectral simulation method for $\ell = N/2$. In that case, the sample is the Fourier transform of (see (16))

$$b'_k := \begin{cases} 0 & k = 0; \\ \sqrt{\frac{f(t_k)}{2N}} \left(U_{k-1}^{(0)} + iU_{k-1}^{(1)} \right) & k = 1, \dots, N/2 - 1; \\ \sqrt{\frac{2f(t_k)}{N}} U_{k-1}^{(0)} & k = N/2; \\ \sqrt{\frac{f(t_k)}{2N}} \left(U_{N-k-1}^{(0)} - iU_{N-k-1}^{(1)} \right) & k = N/2 + 1, \dots, N - 1, \end{cases} \quad (19)$$

where the $U_k^{(i)}$ have their usual meaning. At this point it becomes clear what the relation between the spectral method and the Paxson method is; by comparing b_k to b'_k , we see that the value of $b_{N/2}$ is the only difference (of course, the indexing of the random variables $U^{(i)}$ differs, but this has no impact). However, this effect will vanish for ‘large’ N ; it is readily checked that it does not affect the convergence result of the previous subsection. Thus,

the accuracy of the Paxson method increases with the sample size N , and is even exact as $N \rightarrow \infty$ in the sense that each sample point converges to an ‘exact’ sample point in probability. This new result is the consequence of studying the Paxson method as a special case of the spectral simulation approach.

We can now argue that Paxson’s suggestion to set $\Phi_{N/2} = 0$ is not useful, since this destroys the normality of $b_{N/2}$, which is reflected in *every* sample point, cf. (4). This is the reason that Paxson finds in some cases that the resulting sample points have a ‘nearly’ normal distribution. Therefore, we set

$$b_{N/2} = \sqrt{\frac{f(t_{N/2})}{2N}} U_{N/2}^{(0)},$$

although (as before) the imaginary part has no influence on the sample. In the remaining of this paper, we call this *improved* Paxson method simply the Paxson method.

The method is faster than the Davies and Harte method, although still of order $N \log(N)$. This is because the Paxson method requires just one Fourier transform of a sequence of size N , instead of two transforms of size $2N$. Hence, the Paxson method is approximately four times faster than the exact Davies and Harte method, provided that N is large. It may thus offer a good alternative to exact simulation when large sample sizes are required. Empirical observations show that the Paxson method is slightly less than three times faster for $N = 2^{15}$ when the Paxson approximation for the spectral density (10) is used; apparently the asymptotic regime is not reached yet. A more detailed runtime comparison is made in [5].

The fBm samples produced with Paxson fGn samples have the special property that the end point is always 0. This property is easily seen by using the definition of the FFT transform (4):

$$\sum_{n=0}^{N-1} \sum_{k=0}^{N-1} b_k \exp\left(2\pi i \frac{nk}{N}\right) = \sum_{k=0}^{N-1} b_k \sum_{n=0}^{N-1} \exp\left(2\pi i \frac{nk}{N}\right) = Nb_0 = 0.$$

This highly undesirable property of Paxson fBm samples can be regarded as a fractional Brownian bridge-effect.

Having recognized that the Paxson method equals the spectral simulation method with $\ell = N/2$, a natural question is what can be said about spectral simulation with $\ell = N$. The next subsection shows that there is a connection with the Davies and Harte method.

3.3 The approximate circulant method; connection with the Davies and Harte method

In this subsection, we make the remarkable observation that the spectral simulation approach is related to the circulant diagonalization algorithm of Davies and Harte. In that algorithm, the λ_k are calculated with (4), which can be rewritten as

$$\begin{aligned} \lambda_k &= \sum_{j=0}^{N-1} \gamma(j) \exp\left(2\pi i \frac{jk}{2N}\right) + \sum_{j=1}^{N-1} \gamma(j) \exp\left(2\pi i \frac{(2N-j)k}{2N}\right) \\ &= \sum_{j=-N+1}^{N-1} \gamma(j) \exp\left(2\pi i \frac{jk}{2N}\right). \end{aligned} \tag{20}$$

Letting $t_k = \pi k/N$, we may approximate this by replacing the finite sum by an infinite sum $\sum_{j=-\infty}^{\infty} \gamma(j) \exp(ij t_k)$ when N is ‘large’. Note that this is just the spectral density (6). In that case, we approximate λ_k with $f(t_k)$ for $k = 0, 1, \dots, N$ and with $f(t_{2N-k})$ for $k = N + 1, \dots, 2N - 1$. To avoid problems with the pole, $f(0)$ has to be approximated by a finite value. Since we can compute λ_0 directly from (20) using (2), the choice $f(0) = \lambda_0 = N^{2H} - (N - 1)^{2H}$ seems justified.

Instead of using the exact λ_k , we investigate what happens when the approximations $f(t_k)$ are used in the Davies and Harte algorithm to generate a sample. The FFT is then applied to (see (5))

$$c_k := \begin{cases} \sqrt{\frac{f(t_k)}{2N}} U_k^{(0)} & k = 0; \\ \sqrt{\frac{f(t_k)}{4N}} (U_k^{(0)} + iU_k^{(1)}) & k = 1, \dots, N - 1; \\ \sqrt{\frac{f(t_k)}{2N}} U_k^{(0)} & k = N; \\ \sqrt{\frac{f(t_{2N-k})}{4N}} (U_{2N-k}^{(0)} - iU_{2N-k}^{(1)}) & k = N + 1, \dots, 2N - 1. \end{cases} \quad (21)$$

The first N coefficients of the Fourier transform constitute an approximate fGn sample. Because the only difference with the Davies and Harte method (that was based on a circulant matrix) is the use of *approximate* λ_k , we will refer to this method as *approximate circulant method*. The input coefficients (c_k) closely resemble (16) with $\ell = N$; the differences are the coefficients for $k = 0$ and $k = N$. Again, the effects of this difference vanish as the sample size N grows, which indicates that the method is asymptotically exact in the sense described earlier.

Moreover, we can measure now how the spectral simulation method with $\ell = N$ performs in the spectral domain for finite N : it can be calculated what ‘spectral’ error is made by using the spectral density (6) rather than the λ_k . A further investigation of this issue is found in the next section.

By approximating λ_k with the value of the spectral density at t_k , the number of applications of the FFT on a sequence of size $2N$ is halved. Instead, the spectral density is evaluated in $2N$ points; this is done in order N time. Therefore, the spectral simulation method with $\ell = 2N$ is theoretically twice as fast as the exact Davies and Harte method for large N . However, the method is in practice only a bit faster than the exact method for reasonable sample sizes.

This observation indicates that it makes no sense to increase ℓ further to $2N$ in the spectral simulation method. Still, sometimes (e.g. in time series analysis) the autocovariances of some relevant process are unknown, although the spectral density has some known closed form. Then, it is an interesting option to simulate such a process with the approximate spectral simulation method with $\ell = 2N$. Note that the proof in the appendix indicates that the convergence result holds for any stationary Gaussian process with spectral density $f(\cdot)$.

3.4 Improvements

We have seen that the function $\xi_n^{(\ell)}$ defined in (14) converges to ξ_n in the sense of (18), which led to all theoretical results in the current section. A natural question is if this simple function is the only possible choice for which this type of convergence holds. A glance at the proof indicates that the answer is negative. The function $\xi_n^{(\ell)}$ was constructed by evaluating the spectral density in the rightmost point of each interval of the grid. Looking at (17),

another interesting possibility is to replace $f(t_{k+1})/\ell$ by

$$\int_{t_k}^{t_{k+1}} \frac{f(\lambda)}{\pi} d\lambda. \quad (22)$$

Unfortunately, it depends on the spectral density $f(\cdot)$ whether this integral can be computed (efficiently). In the fGn case, this integral cannot be computed directly, but should be approximated (e.g., by integrating f numerically). For $k = 0$, the integral can be approximated accurately by replacing $f(\lambda)$ by its Taylor expansion $\sin(\pi H)\Gamma(2H + 1)\lambda^{1-2H}$, cf. (8).

Instead of using computationally intensive numerical integration techniques to calculate (22), it is also possible to interpolate $f(\cdot)$ linearly between the grid points (this is impossible for $k = 0$ when $H > 1/2$ because of the pole at 0, but the Taylor based approximation can then be used). This variant is similar to the approach described in Section 3.1; the spectral density is evaluated at the middle point of each interval on the grid, instead of the rightmost point.

These possibilities will not be investigated in the remaining part of this paper; we will focus on the Paxson and the approximate circulant method.

4 Evaluation of the methods

4.1 Error analysis

As already pointed out in the preceding section, the connection between the Davies and Harte method and the approximate circulant method is useful to perform an error analysis.

Define the ‘spectral’ error σ_k^2 for $k > 0$ by

$$\sigma_k^2 := \left(\sqrt{\sum_{j=-\infty}^{\infty} \gamma(j) \exp(\pi i j k / N)} - \sqrt{\sum_{j=-N+1}^{N-1} \gamma(j) \exp(\pi i j k / N)} \right)^2. \quad (23)$$

For this moment, we will neglect the notational technicalities involving the factor $\sqrt{2N}$ instead of $\sqrt{4N}$ in w_N and c_N . Then, σ_k^2 is the variance of $w_k - c_k$, up to a constant scaling factor, see (5) and (21).

Denote the approximate circulant sample by \check{X} and the exact Davies and Harte sample by X . It is readily checked that the ‘time domain’ error $\check{X}_n - X_n$ is a centered Gaussian variable with variance

$$\sigma_{\check{X}_n - X_n}^2 = \frac{1}{2N} \sum_{k=0}^{2N-1} \sigma_k^2,$$

which does not depend on n . The convergence result of the previous section guarantees that $\sigma_{\check{X}_n - X_n}^2$ decays to 0 as $N \rightarrow \infty$.

Notice that this observation concerns the fGn sample, but that it is also interesting to calculate the variance $\sigma_{\check{Y}_n - Y_n}^2$ of the error $\check{Y}_n - Y_n = \sum_{j=0}^n (\check{X}_j - X_j)$ of the corresponding fBm sample. It may be checked that $\sigma_{\check{Y}_n - Y_n}^2$ depends on n and is given by

$$\sigma_{\check{Y}_n - Y_n}^2 = \frac{1}{2N} \sum_{k=0}^{2N-1} \sigma_k^2 \left[\left(\sum_{j=0}^n \cos^2(\pi j k / N) \right)^2 + \left(\sum_{j=0}^n \sin^2(\pi j k / N) \right)^2 \right].$$

Since σ_k^2 is, in the presence of a pole in the spectral density ($H > 1/2$), large when k is either small or large, the term between the square brackets is approximately n^2 for these values of k . This implies that the expected square error of the fBm sample $\sigma_{Y_n - Y_n}^2$ is proportional to n^2 for large sample sizes, i.e. $\sigma_{Y_n - Y_n}^2 \approx Cn^2\sigma_{X_n - X_n}^2$ for some positive constant C . Since $\sigma_{X_n - X_n}^2$ converges to zero as $N \rightarrow \infty$, the expected square error of a fBm sample point at *fixed* n converges to zero as well.

4.2 Numerical analysis of the autocovariance function

Since the autocovariance function characterizes a stationary Gaussian sample uniquely, insight into the performance of the spectral simulation method can be gained by comparing the approximate covariances to the desired covariances.

We have already seen that spectrally simulated fGn samples are stationary. This holds also for Paxson and approximate circulant samples, for which the autocovariance functions $\tilde{\gamma}_P$ resp. $\tilde{\gamma}_{AC}$ are given by

$$\tilde{\gamma}_P(k) = \sum_{j=1}^{N/2-1} \frac{f(2\pi j/N)}{N/2} \cos(2\pi jk/N) + \frac{f(\pi)}{2N} (-1)^k \quad (24)$$

and

$$\tilde{\gamma}_{AC}(k) = \sum_{j=1}^{N-1} \frac{f(\pi j/N)}{N} \cos(\pi jk/N) + \frac{N^{2H} - (N-1)^{2H}}{2N} + \frac{f(\pi)}{2N} (-1)^k \quad (25)$$

for $k = 0, \dots, N-1$.

Some problems arise for the Paxson method at this point. Recall that in a sample of size N , the (co)variances $\gamma(0), \dots, \gamma(N-1)$ are approximated by $\tilde{\gamma}_P(0), \dots, \tilde{\gamma}_P(N-1)$. From (2) follows that the autocovariance function $\gamma(k)$ decreases in k for $H > 1/2$, but this is not the case for the approximation $\tilde{\gamma}_P(k)$, since this function is symmetrical around $k = N/2$.

To illustrate this problem, the autocovariance function $\tilde{\gamma}_P$ is plotted together with the autocovariance function γ for $H = 0.8$ and $N = 2^8$ in the upper panel of Figure 1. The lower panel consists of the functions $\tilde{\gamma}_{AC}$ and γ .

Besides the symmetry problem of $\tilde{\gamma}_P$, the differences with the exact autocovariance function are relatively large; even *negative* covariances are present in Paxson samples for $N = 2^8$. Still, the Paxson method passed many tests (see Paxson [12]), even without the normality improvement of Section 3. It should be noticed that $\tilde{\gamma}_P$ approximates γ better as the sample size N increases, which is probably the reason for Paxson's empirical finding that satisfactory samples are obtained for $N = 2^{15}$.

According to the lower panel of Figure 1, the approximate circulant method is more promising. Except for high-lag autocovariances, the function $\tilde{\gamma}_{AC}$ is almost indistinguishable from γ , which will even improve when $N \rightarrow \infty$.

However, we are particularly interested in the tail behavior of $\tilde{\gamma}_{AC}$ when $H > 1/2$, which cannot be observed in Figure 1. In view of (3), this tail behavior is best studied on a log-log scale; for fGn with Hurst parameter $H = 0.8$, the graph of γ on a log-log scale is a straight line with slope $2H - 2 = -0.4$ for large k . This plot is given in Figure 2 for the approximate circulant method with $H = 0.8$ for $N = 2^{11}$ and $N = 2^{13}$. Since we are mainly interested in the tail behavior, the plot of the autocovariance function starts at $k = 2^{10}$. We see that the exact autocovariance function is indeed a straight line, but that this is not the case for the autocovariance function of the approximate circulant method. The differences at the left end

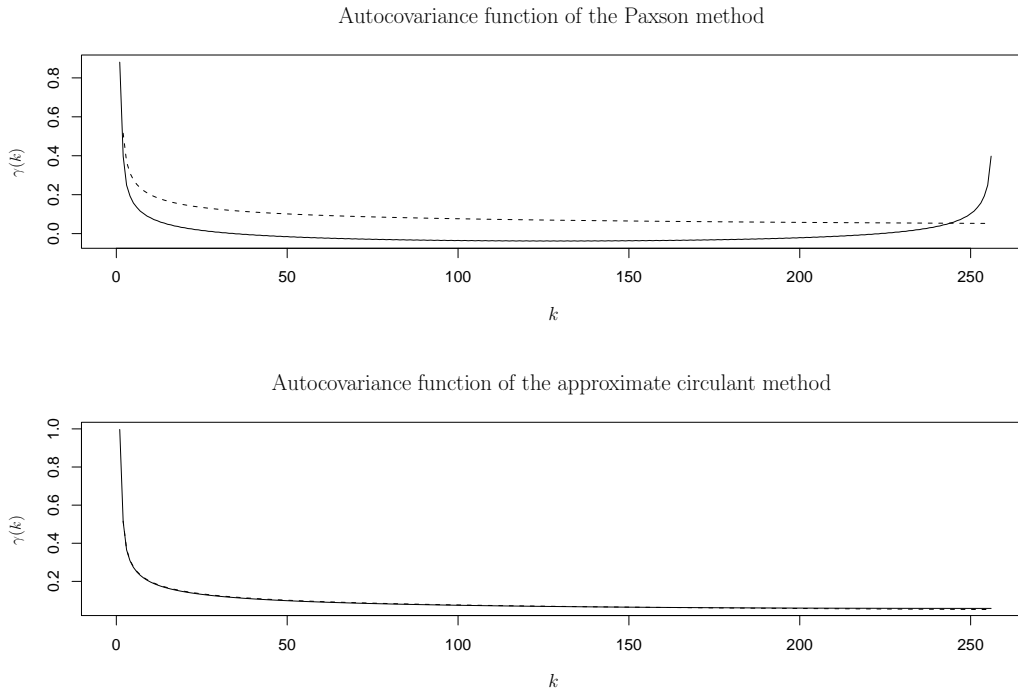


Figure 1: The autocovariance function of the Paxson method and the approximate circulant method with the autocovariance function of fractional Gaussian noise (dashed) for $H = 0.8$.




Figure 2: The tail of the autocovariance function of the approximate circulant method on a log-log scale. The dashed line is the exact autocovariance function ($H = 0.80$).

of the graph are quite small, and can presumably even be made smaller by implementing one of the suggested improvements of Section 3. Interestingly, the solid line curves upwards for the approximate circulant method. Of course, this has to do with the fact that a long-range dependent process is approximated by a short-range dependent process.

If a sample of size N is needed, an obvious improvement is to take the first N elements of an approximate circulant sample of size $2N$. This corresponds to the spectral simulation method with $\ell = 2N$.

5 End notes and conclusions

As we have seen, a convenient property of spectral simulation is the convergence in probability of the approximate fGn sample points to exact fGn sample points. The rate at which this convergence takes place and the implications for the corresponding fBm sample are interesting yet unsolved problems.

The spectral simulation approach for one-dimensional fBm can be extended to higher dimensions (for appropriately defined higher dimensional fBm), for instance by the so-called turning bands method [17]. It should be noted that the Davies and Harte method can also be extended to higher dimensions [6, 16].

The main results can be summarized as follows:

- The described spectral simulation method provides the theoretical foundation for the Paxson method, which clarifies many empirical observations and leads to an important improvement of the method.
- A natural extension of the Paxson method, the approximate circulant method, has a direct connection with the exact Davies and Harte method. The obtained insights show that the accuracy of both methods increases as the sample size grows, and that the methods are even asymptotically exact in the sense that the expected square error at each sample point converges to zero.
- It is possible to perform an error analysis for the approximate circulant method. This shows that the expected square error of the corresponding *fBm* sample at sample point n grows approximately like n^2 .
- Whereas the errors in the covariances are quite large for the Paxson method (although they vanish asymptotically), the autocovariance function of the approximate circulant method is visually almost indistinguishable from the exact autocovariance function. Still, this autocovariance function does *not* show the desired hyperbolic tail behavior when $H > 1/2$. In fact, it is theoretically impossible to obtain this structure in spectrally simulated samples, since spectral density of fGn is approximated by functions without pole at 0.

Appendix: Proof of (18)

We keep n fixed throughout the proof, and start by showing that $\theta_n^{(\ell)}$ converges to θ_n in L^2 -norm as $\ell \rightarrow \infty$, where $\theta_n(\lambda) = \sqrt{\frac{f(\lambda)}{2\pi}} \sin(n\lambda)$. Note that f has a pole at 0, but that the sine-term compensates this pole. Thus, $\theta_n(\lambda)$ is continuous on $[0, \pi]$ for every Hurst parameter $0 < H < 1$. It follows that $[\theta_n(\lambda)]^2$ is Riemann integrable on $[0, \pi]$, which implies

that $\lim_{\ell \rightarrow \infty} \int_0^\pi [\theta_n^{(\ell)}(\lambda)]^2 d\lambda = \int_0^\pi [\theta_n(\lambda)]^2 d\lambda$. Since $\theta_n^{(\ell)}(\lambda) \rightarrow \theta_n(\lambda)$ for every $0 \leq \lambda \leq \pi$ as $\ell \rightarrow \infty$, we have the desired convergence of $\theta_n^{(\ell)}$ to θ_n in L^2 -norm, see for instance Theorem 4.5.4 of Chung [2].

More care is needed to prove a similar fact for the ξ_n functions. When $0 < H \leq 1/2$, the same reasoning as above applies. However, ξ_n has a pole at 0 for $1/2 < H < 1$. Still, f is Riemann integrable because f is a spectral density; in fact, $\int_0^\pi f(\lambda) d\lambda = \pi$. For any $\epsilon > 0$, this makes it possible to find a $\delta > 0$ (independent of ℓ) such that

$$\begin{aligned} \int_0^\delta [\xi_n^{(\ell)}(\lambda) - \xi_n(\lambda)]^2 d\lambda &\leq \int_0^\delta [\xi_n^{(\ell)}(\lambda) + \xi_n(\lambda)]^2 d\lambda \leq 4 \int_0^\delta [\xi_n(\lambda)]^2 d\lambda \\ &\leq 4 \int_0^\delta f(\lambda) d\lambda < \epsilon/2, \end{aligned}$$

where the first two inequalities use the fact that $\xi_n(\lambda) \geq \xi_n^{(\ell)}(\lambda) \geq 0$ for small λ . As before, we have $\xi_n^{(\ell)}(\lambda) \rightarrow \xi_n(\lambda)$ for every $\delta \leq \lambda \leq \pi$, and the Riemann integrability of $[\xi_n(\lambda)]^2$ on $[\delta, \pi]$, which implies that $\lim_{\ell \rightarrow \infty} \int_\delta^\pi [\xi_n^{(\ell)}(\lambda)]^2 d\lambda = \int_\delta^\pi [\xi_n(\lambda)]^2 d\lambda$. Hence, it is possible to find an ℓ with the property that $\int_\delta^\pi [\xi_n^{(\ell)}(\lambda) - \xi_n(\lambda)]^2 d\lambda < \epsilon/2$ and the claim is proven.

Note that the arguments of the proof apply for a general spectral density, not necessarily the spectral density of fGn.

References

- [1] S. ASMUSSEN, *Stochastic simulation with a view towards stochastic processes*. Notes of a Concentrated Advance Course at MaPhySto, Aarhus, Denmark, February 1999.
- [2] K. L. CHUNG, *A Course in Probability Theory*, Academic Press, San Diego, third ed., 2001.
- [3] M. S. CROUSE AND R. G. BARANIUK, *Fast, exact synthesis of Gaussian and nonGaussian long-range dependent processes*. Submitted to *IEEE Transactions on Information Theory*, 1999.
- [4] R. B. DAVIES AND D. S. HARTE, *Tests for Hurst effect*, *Biometrika*, 74 (1987), pp. 95–102.
- [5] A. B. DIEKER, *Simulation of fractional Brownian motion*, Master's thesis, Vrije Universiteit Amsterdam, April 2002. See <http://www.cwi.nl/~ton>.
- [6] C. R. DIETRICH AND G. N. NEWSAM, *Fast and exact simulation of stationary Gaussian processes through circulant embedding of the covariance matrix*, *SIAM Journal on Scientific Computing*, 18 (1997), pp. 1088–1107.
- [7] J. R. M. HOSKING, *Modeling persistence in hydrological time series using fractional differencing*, *Water Resources Research*, 20 (1984), pp. 1898–1908.
- [8] I. KARATZAS AND S. E. SHREVE, *Brownian motion and stochastic calculus*, Springer-Verlag, New York, second ed., 1991.

- [9] G. LINDGREN, *Lectures on stationary stochastic processes*, Lund University, 1999. See <http://www.maths.lth.se/matstat/staff/georg>.
- [10] B. B. MANDELBROT AND J. W. VAN NESS, *Fractional Brownian motions, fractional noises and applications*, SIAM Review, 10 (1968), pp. 422–437.
- [11] I. NORROS, P. MANNERSALO, AND J. L. WANG, *Simulation of fractional Brownian motion with conditionalized random midpoint displacement*, Advances in Performance Analysis, 2 (1999), pp. 77–101.
- [12] V. PAXSON, *Fast, approximate synthesis of fractional Gaussian noise for generating self-similar network traffic*, Computer Communication Review, 27 (1997), pp. 5–18.
- [13] M. B. PRIESTLEY, *Spectral analysis and time series*, vol. 1, Academic Press, 1981.
- [14] D. SAUPE, *Algorithms for random fractals*, in The Science of Fractal Images, H.-O. Peitgen and D. Saupe, eds., Springer-Verlag, 1988.
- [15] Y. G. SINAI, *Self-similar probability distributions*, Theory of Probability and its Applications, 21 (1976), pp. 64–80.
- [16] A. T. A. WOOD AND G. CHAN, *Simulation of stationary Gaussian processes in $[0, 1]^d$* , Journal of Computational and Graphical Statistics, 3 (1994), pp. 409–432.
- [17] Z.-M. YIN, *New methods for simulation of fractional Brownian motion*, Journal of Computational Physics, 127 (1996), pp. 66–72.