

A COMPARATIVE STUDY OF THREE SELF-SIMILAR TELETRAFFIC GENERATORS

Hae-Duck J. Jeong[†], Don McNickle[‡] and Krzysztof Pawlikowski[†]

Department of [†]Computer Science and [‡]Management
University of Canterbury
Christchurch
New Zealand

E-mail: {joshua, krysz}@cosc.canterbury.ac.nz

KEYWORDS

Teletraffic Generators, Complexity, Self-Similar Processes, Hurst Parameter.

ABSTRACT

It is generally accepted that *self-similar* (or *fractal*) processes may provide better models for teletraffic in modern telecommunication networks than Poisson processes. If this is not taken into account, it can lead to inaccurate conclusions about performance of telecommunication networks. Thus, an important requirement for conducting simulation studies of telecommunication networks is the ability to generate long synthetic stochastic self-similar sequences.

Three generators of pseudo-random self-similar sequences, based on the FFT (Paxson 1997), RMD (Lau et al. 1995) and SRA method (Crilly et al. 1991; Jeong et al. 1998) are compared and analysed in this paper. Properties of these generators were experimentally studied in the sense of their statistical accuracy and times required to produce sequences of a given (long) length. While all three generators show similar levels of accuracy of the output data (in the sense of relative accuracy of the Hurst parameter), the RMD- and SRA-based generators appear to be much faster than the generator based on FFT. Our results also show that a robust method for comparative studies of self-similarity in pseudo-random sequences is needed.

INTRODUCTION

The search for accurate mathematical models of data streams in modern telecommunication networks has attracted a considerable amount of interest in the last few years. The reason is that several recent teletraffic studies of local and wide area networks, including the world wide web, have shown that commonly used teletraffic models, based on Poisson or related processes, are not able to capture the self-similar (or fractal) nature of teletraffic (Leland et al. 1994; Likhanov

et al. 1995; Paxson and Floyd 1995; Ryu 1996), especially when they are engaged in such sophisticated services as variable-bit-rate (VBR) video transmission (Garrett and Willinger 1994; Krunz and Makowski 1998; Rose 1997). The properties of teletraffic in such scenarios are very different from both the properties of conventional models of telephone traffic and the traditional models of data traffic generated by computers.

The use of traditional models of teletraffic can result in overly optimistic estimates of performance of telecommunication networks, insufficient allocation of communication and data processing resources, and difficulties in ensuring the quality of service expected by network users (Beran 1992; Neidhardt and Wang 1998; Paxson and Floyd 1995). On the other hand, if the strongly correlated character of teletraffic is explicitly taken into account, this can also lead to more efficient traffic control mechanisms.

Several methods for generating pseudo-random self-similar sequences have been proposed. They include methods based on fast fractional Gaussian noise (Mandelbrot 1971), fractional ARIMA processes (Hosking 1984), the $M/G/\infty$ queue model (Krunz and Makowski 1998; Leland et al. 1994), autoregressive processes (Cario and Nelson 1998; Granger 1980), spatial renewal processes (Taralp et al. 1998), etc. Some of them generate asymptotically self-similar sequences and require large amounts of CPU time. For example, Hosking's method (Hosking 1984), based on the F-ARIMA(0, d , 0) process, needs many hours to produce a self-similar sequence with 131,072 (2^{17}) numbers on a Sun SPARCstation 4 (Leland et al. 1994). It requires $O(n^2)$ computations to generate n numbers. Even though exact methods of generation of self-similar sequences exist (for example: (Mandelbrot 1971)), they are only fast enough for short sequences. They are usually inappropriate for generating long sequences because they require multiple passes along generated sequences. To overcome this, approximate methods for generation of self-similar sequences in simulation studies of telecommunication networks have also been proposed (Lau et al. 1995; Paxson 1997).

Our comparative evaluation of three methods proposed for generating self-similar sequences concentrates on two aspects: (i) how accurately self-similar processes can be generated, and (ii) how fast the methods generate long self-similar sequences. We consider three methods: (i) a method based on the *fast Fourier transform (FFT)* algorithm and implemented by Paxson (Paxson 1997); (ii) a method based on the *random midpoint displacement (RMD)* algorithm and implemented by Lau, Erramilli, Wang and Willinger (Lau et al. 1995); and (iii) a method based on the *successive random addition (SRA)* algorithm, proposed by Saupe, D. (Crilly et al. 1991) and implemented by Jeong, McNickle and Pawlikowski (Jeong et al. 1998).

A summary of the basic properties of self-similar processes is given in section . In section the three generators of pseudo-random self-similar sequences are described. Numerical results of comparative analysis of sequences generated by these generators are discussed in section .

SELF-SIMILAR PROCESSES AND THEIR PROPERTIES

Basic definitions of self-similar processes are as follows: A continuous-time stochastic process $\{X_t\}$ is strongly *self-similar* with a self-similarity parameter H ($0 < H < 1$), known as the Hurst parameter, if for any positive stretching factor c , the rescaled process with time scale ct , $c^{-H}X_{ct}$, is equal in distribution to the original process $\{X_t\}$ (Beran 1994). This means that, for any sequence of time points t_1, t_2, \dots, t_n , and for all $c > 0$, $\{c^{-H}X_{ct_1}, c^{-H}X_{ct_2}, \dots, c^{-H}X_{ct_n}\}$ has the same distribution as $\{X_{t_1}, X_{t_2}, \dots, X_{t_n}\}$.

In discrete-time case, let $\{X_k\} = \{X_k : k = 0, 1, 2, \dots\}$ be a (discrete-time) stationary process with mean μ , variance σ^2 , and autocorrelation function (ACF) $\{\rho_k\}$, for $k = 0, 1, 2, \dots$, and let $\{X_k^{(m)}\}_{k=1}^{\infty} = \{X_1^{(m)}, X_2^{(m)}, \dots\}$, $m = 1, 2, 3, \dots$, be a sequence of batch means, i.e., $X_k^{(m)} = (X_{km-m+1} + \dots + X_{km})/m$, $k \geq 1$.

The process $\{X_k\}$ with $\rho_k \rightarrow k^{-\beta}$, as $k \rightarrow \infty$, $0 < \beta < 1$, is called *exactly self-similar* with $H = 1 - (\beta/2)$, if $\rho_k^{(m)} = \rho_k$, for any $m = 1, 2, 3, \dots$. In other words, the process $\{X_k\}$ and the averaged processes $\{X_k^{(m)}\}$, $m \geq 1$, have identical correlation structure.

The process $\{X_k\}$ is *asymptotically self-similar* with $H = 1 - (\beta/2)$, if $\rho_k^{(m)} \rightarrow \rho_k$, as $m \rightarrow \infty$.

The most frequently studied models of self-similar traffic belong either to the class of fractional autoregressive integrated moving-average (F-ARIMA) processes or to the class of fractional Gaussian noise processes; see (Hosking 1984; Leland et al. 1994; Paxson 1997). F-ARIMA(p, d, q) processes were introduced by Hosking (Hosking 1984) who showed that they are asymptotically self-similar with Hurst

parameter $H = d + \frac{1}{2}$, as long as $0 < d < \frac{1}{2}$. In addition, the incremental process $\{Y_k\} = \{X_k - X_{k-1}\}$, $k \geq 0$, is called the *fractional Gaussian noise (FGN)* process, where $\{X_k\}$ designates a fractional Brownian motion (FBM) random process. This process is a (discrete-time) stationary Gaussian process with mean μ , variance σ^2 and $\{\rho_k\} = \{\frac{1}{2}(|k+1|^{2H} - 2|k|^{2H} + |k-1|^{2H})\}$, $k > 0$. A FBM process, which is the sum of FGN increments, is characterised by three properties (Mandelbrot and Wallis 1969): (i) it is a continuous zero-mean Gaussian process $\{X_t\} = \{X_s : s \geq 0 \text{ and } 0 < H < 1\}$ with ACF given by $\rho_{s,t} = \frac{1}{2}(s^{2H} + t^{2H} - |s-t|^{2H})$ where s is time lag and t is time; (ii) its increments $\{X_t - X_{t-1}\}$ form a stationary random process; (iii) it is self-similar with Hurst parameter H , that is, for all $c > 0$, $\{X_{ct}\} = \{c^H X_t\}$, in the sense that, if time is changed by the ratio c , then $\{X_t\}$ is changed by c^H .

Main properties of self-similar processes include (Beran 1994; Cox 1984; Leland et al. 1994):

- *Slowly decaying variance.* The variance of the sample mean decreases more slowly than the reciprocal of the sample size, that is, $Var[\{X_k^{(m)}\}] \rightarrow c_1 m^{-\beta_1}$ as $m \rightarrow \infty$, where c_1 is a constant and $0 < \beta_1 < 1$.
- *Long-range dependence.* A process $\{X_k\}$ is called a stationary process with *long-range dependence (LRD)* if its ACF $\{\rho_k\}$ is non-summable, that is, $\sum_{k=0}^{\infty} \rho_k = \infty$. The speed of decay of autocorrelations is more like hyperbolic than exponential.
- *Hurst effect.* Self-similarity manifests itself by a straight line of slope β_2 on a log-log plot of the R/S statistic. For a given set of numbers $\{X_1, X_2, \dots, X_n\}$ with sample mean $\hat{\mu} = E\{X_i\}$ and sample variance $S^2(n) = E\{(X_i - \hat{\mu})^2\}$, Hurst parameter H is presented by the *rescaled adjusted range* $\frac{R(n)}{S(n)}$ (or R/S statistic) where $R(n) = \max\{\sum_{i=1}^k (X_i - \hat{\mu}), 1 \leq k \leq n\} - \min\{\sum_{i=1}^k (X_i - \hat{\mu}), 1 \leq k \leq n\}$ and S is estimated by $S(n) = \sqrt{E\{(X_i - \hat{\mu})^2\}}$. Hurst found empirically that for many time series observed in nature the expected value of $\frac{R(n)}{S(n)}$ asymptotically satisfies the power law relation, i.e., $E[\frac{R(n)}{S(n)}] \rightarrow c_2 n^H$ as $n \rightarrow \infty$ with $0.5 < H < 1$ and c_2 is a finite positive constant (Beran 1994).

In simulation of telecommunication networks, given a sequence of the approximate FBM process $\{X_t\}$, we can obtain a self-similar cumulative arrival process $\{Y_t\}$ (Lau et al. 1995; Norros 1994): $\{Y_t\} = Mt + \sqrt{AM}\{X_t\}$, $t \in (-\infty, +\infty)$ where M is the mean input rate and A is the peakedness factor, defined as the ratio of variance to the mean, $M > 0, A > 0$. The Gaussian incremental process $\{\tilde{Y}_t\}$ from time t to time $t+1$ is given as: $\{\tilde{Y}_t\} = M + \sqrt{AM}[\{X_{t+1}\} - \{X_t\}]$.

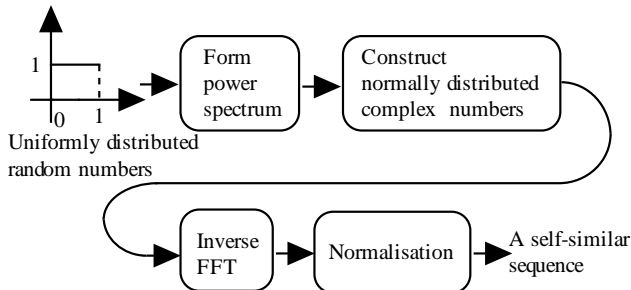


Figure 1: FFT method




Figure 3: The first three steps in the RMD method

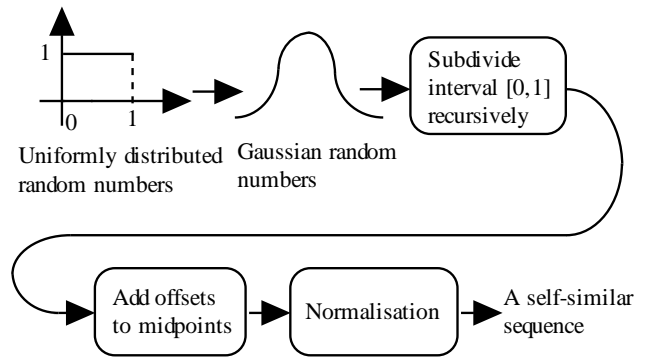


Figure 2: RMD method

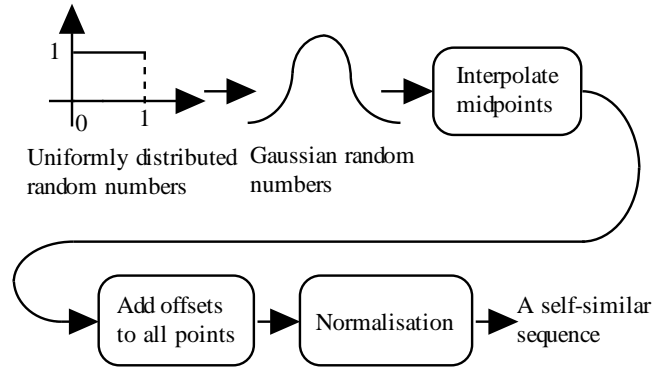


Figure 4: SRA method

THREE METHODS

The FFT- and RMD-based methods were suggested as being sufficiently fast for practical applications in generation of simulation input data (Lau et al. 1995; Paxson 1997). In this paper, we have reported properties of these two methods and compare them with SRA, one of recently proposed alternative methods for generating pseudo-random self-similar sequences (Jeong et al. 1998). These methods can be characterised as follows:

FFT Method

This method generates approximate self-similar sequences based on the Fast Fourier Transform and a process known as the Fractional Gaussian Noise (FGN) process, (Figure 1.) Its main difficulty is connected with calculating the power spectrum, which involves an infinite summation. Paxson has solved this problem by applying a special approximation.

Figure 1 shows how the FFT method generates self-similar sequences. Briefly, it is based on (i) calculation of the power spectrum using the periodogram (the power spectrum at a given frequency represents an independent exponential random variable); (ii) construction of complex numbers which are governed by the normal distribution; (iii) execution of the inverse FFT. For a more detailed reference, see (Paxson

1997).

RMD Method

The basic concept of the *random midpoint displacement (RMD)* algorithm is to extend the generated sequence recursively, by adding new values at the midpoints from the values at the endpoints.

Figure 2 outlines how the RMD algorithm works. Figure 3 illustrates the first three steps of the method, leading to generation of the sequence $(d_{3,1}, d_{3,2}, d_{3,3}, d_{3,4})$. The reason for subdividing the interval between 0 and 1 is to construct the Gaussian increments of X . Adding offsets to midpoints makes the marginal distribution of the final result normal. For more detailed discussions of the RMD method, see (Lau et al. 1995; Peitgen et al. 1992).

SRA Method

Another alternative method for the direct generation of FBM process is based on the *successive random addition (SRA)* algorithm (Crilly et al. 1991). The SRA method uses the midpoints like RMD, but adds a displacement of a suitable variance to all of the points to increase stability of the generated sequence (Peitgen and Saupe 1988).

Figure 4 shows how the SRA method generates an approximate self-similar sequence. The reason for interpolating midpoints is to construct Gaussian increments of X , which are correlated. Adding offsets to all points should make the resulted sequence self-similar and of normal distribution (Peitgen and Saupe 1988).

The SRA method consists of the following steps:

- Step.1 If the process $\{X_t\}$ is to be computed for times instances t between 0 and 1, then start out by setting $X_0 = 0$ and selecting X_1 as a pseudo-random number from a Gaussian distribution with mean 0 and variance $Var[X_1] = \sigma_0^2$. Then $Var[X_1 - X_0] = \sigma_0^2$.
- Step.2 Next, $X_{\frac{1}{2}}$ is constructed by the interpolation of the midpoint, that is, $X_{\frac{1}{2}} = \frac{1}{2}(X_0 + X_1)$.
- Step.3 Add a displacement of a suitable variance to all of the points, i.e., $X_0 = X_0 + d_{1,1}$, $X_{\frac{1}{2}} = X_{\frac{1}{2}} + d_{1,2}$, $X_1 = X_1 + d_{1,3}$. The offsets $d_{1,*}$ are governed by fractional Gaussian noise. For $Var[X_{t_2} - X_{t_1}] = |t_2 - t_1|^{2H} \sigma_0^2$ to be true, for any $t_1, t_2, 0 \leq t_1 \leq t_2 \leq 1$, it is required that $Var[X_{\frac{1}{2}} - X_0] = \frac{1}{4}Var[X_1 - X_0] + 2S_1^2 = (\frac{1}{2})^{2H} \sigma_0^2$, that is, $S_1^2 = \frac{1}{2}(\frac{1}{2})^{2H} (1 - 2^{2H-2}) \sigma_0^2$.
- Step.4 Next, Step.2 and Step.3 are repeated. Therefore, $S_n^2 = \frac{1}{2}(\frac{1}{2^n})^{2H} (1 - 2^{2H-2}) \sigma_0^2$, where σ_0^2 is an initial variance and $0 < H < 1$.

Using the above steps, the SRA method generates an approximate self-similar FBM process.

ANALYSIS OF SELF-SIMILAR SEQUENCES

Three generators of self-similar sequences of pseudo-random numbers described in the Section have been implemented in C on a Sun SPARCstation 4 (110 MHz, 32 MB), and used to generate self-similar cumulative arrival processes, mentioned at the end of Section . The mean times required for generating sequences of a given length were obtained by using the SunOS 5.5 date command and averaged over 30 iterations, having generated sequences of 32,768 (2^{15}), 131,072 (2^{17}), 262,144 (2^{18}), 524,288 (2^{19}) and 1,048,576 (2^{20}) numbers.

We have also analysed the efficiency of these methods in the sense of their accuracy. For each of $H = 0.5, 0.55, 0.7, 0.9, 0.95$, each method was used to generate over 100 sample sequences of 32,768 (2^{15}) numbers starting from different random seeds. Self-similarity and marginal distributions of the generated sequences were assessed by applying the best currently available techniques. These include:

- *Anderson-Darling goodness-of-fit test*: used to show that the marginal distribution of sample sequences generated by all three methods is normal or almost normal, since all three methods are based on Gaussian

processes. This test is more powerful than *Kolmogorov-Smirnov* when testing against a specified normal distribution (Gibbons and Chakraborti 1992).

- *Periodogram plot*: used to show whether a generated sequence is LRD or not. It can be shown that if the autocorrelations were summable, then near the origin the periodogram should be scattered randomly around a constant. If the autocorrelations were non-summable, i.e., LRD, the points of a sequence are scattered around a negative slope. The periodogram plot is obtained by plotting $\log_{10}(\text{periodogram})$ against $\log_{10}(\text{frequency})$. An estimate of the Hurst parameter is given by $\hat{H} = (1 - \hat{\beta}_3)/2$ where $\hat{\beta}_3$ is the slope (Beran 1994).
- *R/S statistic plot*: graphical R/S analysis of empirical data can be used to estimate the Hurst parameter \hat{H} . An estimate of H is given by the asymptotic slope $\hat{\beta}_2$ of the R/S statistic plot, i.e., $\hat{H} = \hat{\beta}_2$ (Beran 1994).
- *Variance-time plot*: is obtained by plotting $\log_{10}(Var(X^{(m)}))$ against $\log_{10}(m)$ and by fitting a simple least square line through the resulting points in the plane. An estimate of the Hurst parameter is given by $\hat{H} = 1 - \hat{\beta}_1/2$ where $\hat{\beta}_1$ is the slope (Beran 1994).
- *Whittle's approximate maximum likelihood estimate (MLE)*: is a more refined data analysis method to obtain confidence intervals (CIs) for the Hurst parameter H (Beran 1994).

Analysis of Accuracy

We have summarised the results of our analysis in the following:

- Anderson-Darling goodness-of-fit test was applied to test normality of sample sequences. The results of the tests, executed at the 5% significance level, showed that the generated sequences could be considered as normally distributed for all but a few sequences with the high value of H .

The estimates of Hurst parameter obtained from the periodogram, the R/S statistic, the variance-time and Whittle's MLE, have been used to compare the accuracy of the three methods. The relative inaccuracy ΔH is calculated using the formula: $\Delta H = \frac{\hat{H} - H}{H} * 100\%$, where H is the input value and \hat{H} is an empirical mean value. The presented numerical results are all averaged over 100 sequences.

- The periodogram plots have slopes decreasing as H increases. The negative slopes of all our plots for $H = 0.5, 0.55, 0.7, 0.9, 0.95$ were the evidence of self-similarity. A comparison of relative inaccuracy ΔH of the estimated Hurst parameters of three methods using periodogram plot is given in Table 1. We see that in the most cases parameter H of the FFT method was closer to the required value than in the case of the RMD

and SRA methods, although the relative inaccuracy degrades with increasing H (but never exceeds 6%). The analysis of periodograms suggest that the FFT method always produces self-similar sequences with positively biased H , while sequences produced by two other methods are negatively biased.

- The plots of R/S statistic clearly confirmed the self-similar nature of the generated sequences. The relative inaccuracy ΔH of the estimated Hurst parameter, obtained by R/S statistic plot, is given in Table 2. As we see, these results suggest that the FFT method is slightly better than the other two (but for $H = 0.9, 0.95$). This method of analysis of H does not link any of these generators with persistently negative or positive bias of H , as the periodogram plots did.
- The variance-time plots also supported the claim that generated sequences were self-similar. Table 3 gives the relative inaccuracy ΔH of the estimated Hurst parameters obtained by the variance-time plot. Again, all three methods show comparable quality of the output sequences in the sense of H , with the relative inaccuracy increasing with the increase in H , but remaining below 8%. This time, all results but one suggest that the output sequences are negatively biased H , regardless of the method.
- The results for Whittle estimator of H with the corresponding 95% CIs $\hat{H} \pm 1.96\hat{\sigma}_{\hat{H}}$, see Table 4, show that for all input H values, CIs for the FFT method cover the assumed theoretical values, while the RMD and SRA methods produce sequences weaker correlated than expected (except $H = 0.5$).

Our results show that all three generators produce approximately self-similar sequences, with the relative inaccuracy ΔH increasing with H , but always staying below 9%. Apparently there is a problem with more detailed comparative studies of such generators, since different methods of analysis of the Hurst parameter can give different results regarding the bias of \hat{H} characterising the same output sequences. More reliable methods for assessment of self-similarity in pseudo-random sequences are needed.

Computational Complexity

The results of our experimental analysis of mean times needed by the three generators for generating pseudo-random self-similar sequences of a given length are shown in Table 5. The main conclusions are listed below.

- *FFT method* is the slowest of the three analysed methods. This is caused by relatively high complexity of the inverse FFT algorithm. Table 5 shows its time complexity and the mean running time. It took 5 seconds to generate a sequence of 32,768 (2^{15}) numbers, while generation of a sequence with 1,048,576 (2^{20}) numbers

took 3 minutes and 47 seconds. FFT method requires $O(n \log n)$ computations to generate n numbers (Press et al. 1986).

- *RMD method* is faster and simpler than FFT. Table 5 shows its time complexity and the mean running time. Generation of a sequence with 32,768 (2^{15}) numbers took 3 seconds. It also took 1 minute and 33 seconds to generate a sequence of 1,048,576 (2^{20}) numbers. The theoretical algorithmic complexity is $O(n)$ (Peitgen and Saupe 1988).
- *SRA method* appears to be as fast as RMD. Table 5 shows its time complexity and the mean running time. The theoretical algorithmic complexity is $O(n)$ (Peitgen and Saupe 1988).

In summary, our results show that the generators based on RMD and SRA are faster in practical applications than the generator based on FFT, when long self-similar sequences of numbers are needed.

CONCLUSIONS

In this paper we have presented the results of a comparative analysis of three generators of (long) pseudo-random self-similar sequences. It appears that all three generators, based on FFT, RMD and SRA, generate approximately self-similar sequences, with the relative inaccuracy of the resulted H below 9%, if $0.5 \leq H \leq 0.95$. On the other hand, the analysis of mean times needed for generating sequences of given lengths shows that two generators (based on RMD and SRA) should be recommended for practical simulation of telecommunication networks, since they are much faster than the generator based on FFT. Our study has also revealed that a robust method for comparative studies of self-similarity in pseudo-random sequences is needed, since currently available methods can provide inconclusive proofs of accuracy of such sequences. This is the direction of our current research.

REFERENCES

- Beran, J. 1992. "Statistical Methods for Data with Long Range Dependence." *Statistical Science* 7, no.4: 404-427.
- Beran, J. 1994. *Statistics for Long-Memory Processes*, Chapman and Hall, New York.
- Cario, M.C. and B.L. Nelson. 1998. "Numerical Methods for Fitting and Simulating Autoregressive-to-Anything Processes." *INFORMS Journal on Computing* 10, no.1: 72-81.
- Cox, D.R.1984. "Long-Range Dependence: a Review." *Statistics: An Appraisal*. Iowa State Statistical Library, The Iowa State University Press, 55-74.

Crilly, A.J.; and R.A. Earnshaw; and H. Jones. 1991. *Fractals and Chaos*. Springer-Verlag.

Garrett, M.W. and W. Willinger. 1994. "Analysis, Modeling and Generation of Self-Similar VBR Video Traffic." *Computer Communication Review Proceedings of ACM SIGCOMM'94* (London, UK, Aug.) 24, no. 4: 269-280.

Gibbons, J.D. and S. Chakraborti. 1992. *Nonparametric Statistical Inference*. Marcel Dekker, Inc.

Granger, C.W.J. 1980. "Long Memory Relationships and the Aggregation of Dynamic Models." *Journal of Econometrics* 14: 227-238.

Hosking, J.R.M. 1984. "Modeling Persistence in Hydrological Time Series Using Fractional Differencing." *Water Resources Research* 20, no. 12 (Dec.): 1898-1908.

Jeong, H.-D.J.; D. McNickle; and K. Pawlikowski. 1998. "A Generator of Pseudo-random Self-Similar Sequences Based on SRA." Technical report TR-COSC 9/98. Department of Computer Science, University of Canterbury, Christchurch, New Zealand.

Krunz, M. and A. Makowski. 1998. "A Source Model for VBR Video Traffic Based on $M/G/\infty$ Input Processes." *Proceedings of IEEE INFOCOM'98* (San Francisco, CA, USA, Mar.), 1441-1448.

Lau, W.-C.; A. Erramilli; J.L. Wang; and W. Willinger. 1995. "Self-Similar Traffic Generation: the Random Midpoint Displacement Algorithm and Its Properties." *Proceedings of IEEE ICC'95* (Seattle, WA), 466-472.

Leland, W.E.; M.S. Taqqu; W. Willinger; and D.V. Wilson. 1994. "On the Self-Similar Nature of Ethernet Traffic (Extended Version)." *IEEE/ACM Transactions on Networking* 2, no. 1: 1-15.

Likhanov, N.; B. Tsybakov; and N.D. Georganas. 1995. "Analysis of an ATM Buffer with Self-Similar ("Fractal") Input Traffic." *Proceedings of IEEE INFOCOM'95*, 985-992.

Mandelbrot, B.B. 1971. "A Fast Fractional Gaussian Noise

Generator." *Water Resources Research* 7: 543-553.

Mandelbrot, B.B. and J.R. Wallis. 1969. "Computer Experiments with Fractional Gaussian Noises." *Water Resources Research* 5, no. 1: 228-267.

Neidhardt, A.L. and J.L. Wang. 1998. "The Concept of Relevant Time Scales and Its Application to Queueing Analysis of Self-Similar Traffic (or Is Hurst Naughty or Nice?)." *Proceedings ACM SIGMETRICS'98* (Madison, Wisconsin, USA, Jun.), 222-232.

Norros, I. 1994. "A Storage Model with Self-Similar Input." *Queueing Systems* 16: 387-396.

Paxson, V. 1997. "Fast, Approximate Synthesis of Fractional Gaussian Noise for Generating Self-Similar Network Traffic." *ACM SIGCOMM, Computer Communication Review* 27, no. 5: 5-18.

Paxson, V. and S. Floyd. 1995. "Wide-Area Traffic: the Failure of Poisson Modeling." *IEEE/ACM Transactions on Networking* 3, no. 3 (Jun.): 226-244.

Peitgen, H.-O.; H. Jurgens; and D. Saupe. 1992. *Chaos and Fractals: New Frontiers of Science*. Springer-Verlag.

Peitgen, H.-O. and D. Saupe. 1988. *The Science of Fractal Images*. Springer-Verlag.

Press, W.H.; B.P. Flannery; S.A. Teukolsky; and W.T. Vetterling. 1986. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press.

Rose, O. 1997. *Traffic Modeling of Variable Bit Rate MPEG Video and Its Impacts on ATM Networks*. PhD thesis, Bayerischen Julius-Maximilians-Universität Würzburg.

Ryu, B.K. 1996. *Fractal Network Traffic: from Understanding to Implications*. PhD thesis, Graduate School of Arts and Science, Columbia University.

Taralp, T.; M. Devetsikiotis; and I. Lambadaris. 1998. "Efficient Fractional Gaussian Noise Generation Using the Spatial Renewal Process." *Proceedings IEEE International Communications Conference (ICC'98)* (Atlanta, GA, USA, Jun.), S41-3.1-S41-3.5.

Table 1: Relative inaccuracy ΔH estimated from periodogram plots.

H	FFT	RMD	SRA
0.5	+ 0.07 %	- 0.01 %	- 0.09 %
0.55	+ 1.26 %	- 1.31 %	- 1.41 %
0.7	+ 3.14 %	- 3.74 %	- 3.78 %
0.9	+ 3.93 %	- 5.10 %	- 5.13 %
0.95	+ 3.99 %	- 5.28 %	- 5.31 %

Table 2: Relative inaccuracy ΔH estimated from R/S statistic plots.

H	FFT	RMD	SRA
0.5	+7.34 %	+8.74 %	+8.71 %
0.55	+5.32 %	+6.28 %	+6.23 %
0.7	+0.82 %	+1.28 %	+1.26 %
0.9	- 5.02 %	- 4.46 %	- 4.44 %
0.95	- 6.89 %	- 6.34 %	- 6.31 %

Table 3: Relative inaccuracy ΔH estimated from variance-time plots.

H	FFT	RMD	SRA
0.5	- 0.85 %	+0.57 %	- 2.76 %
0.55	- 1.00 %	- 0.19 %	- 2.97 %
0.7	- 1.88 %	- 1.76 %	- 3.38 %
0.9	- 5.39 %	- 5.29 %	- 6.00 %
0.95	- 6.98 %	- 6.91 %	- 7.47 %

Table 4: Estimated mean values of H using Whittle's MLE. Each CI is for over 100 sample sequences. 95% CIs for the means are given in parentheses.

Method	Theoretical Hurst parameter				
	.5	.55	.7	.9	.95
FFT	.500 (.490, .510)	.550 (.540, .560)	.700 (.691, .710)	.900 (.891, .909)	.949 (.940, .958)
RMD	.500 (.490, .510)	.538 (.528, .548)	.658 (.647, .666)	.826 (.817, .835)	.870 (.861, .879)
SRA	.500 (.490, .510)	.538 (.528, .547)	.656 (.647, .666)	.825 (.816, .834)	.869 (.860, .878)

Table 5: Complexity and mean running times of generators. Running times were obtained by using the SunOS 5.5 date command on a Sun SPARCstation 4 (110 MHz, 32 MB); each mean is averaged over 30 iterations.

Method	Complexity	Sequence of				
		32,768 Numbers	131,072 Numbers	262,144 Numbers	524,288 Numbers	1,048,576 Numbers
		Mean running time (<i>minute:second</i>)				
FFT	$O(n \log n)$	0:5	0:20	0:35	1:12	3:47
RMD	$O(n)$	0:3	0:11	0:29	0:40	1:33
SRA	$O(n)$	0:3	0:10	0:20	0:40	1:31

BIOGRAPHY