

May 3, 2002

To the Graduate School:

This thesis entitled “Simulation of Self-similar Network Traffic Using High Variance ON/OFF Sources” and written by Philip M. Wells is presented to the Graduate School of Clemson University. I recommend that it be accepted in partial fulfillment of the requirements for the degree of Master of Science with a major in Computer Science.

Dr. James M. Westall

We have reviewed this thesis
and recommend its acceptance:

Dr. Alan Wayne Madison

Dr. Mark Smotherman

Accepted for the Graduate School:

SIMULATION OF SELF-SIMILAR NETWORK
TRAFFIC USING HIGH VARIANCE ON/OFF

SOURCES

A Thesis

Presented to
the Graduate School of
Clemson University

In Partial Fulfillment
of the Requirements for the Degree
Master of Science
Computer Science

by

Philip M. Wells

May 2002

Advisor: Dr. James M. Westall

Abstract

Realistic traffic models are a fundamental requirement for understanding network hardware and software design issues such as queuing behavior, congestion management and buffer sizing. It has been determined that many real-world traffic traces are statistically consistent with long-range dependent or self-similar traffic models, however many current self-similar traffic models are mathematically complex and inherently slow, or are highly parameterized and tailored to a specific type of network traffic (e.g. HTTP). Through simulation, we investigate the range and sensitivity of the few parameters necessary for the generation of traffic from an aggregated, ON/OFF-source model whose ON/OFF-period lengths have infinite variance. Using this simple model, which provides insight into the cause of self-similar traffic, we investigate the range of simulation parameter values that produce traffic which is, to varying degrees, self-similar in nature. This understanding will make it easier to incorporate self-similar traffic models into network performance simulations. We also show that traffic can be simulated that is self-similar in nature with as few as eight ON/OFF-sources, greatly increasing the ease with which this method can be used to synthesize traffic on a small network of machines.

Acknowledgements

I would like to thank my advisor, Dr. Mike Westall for his patience, guidance, and attention to details that I would easily overlook. I would also like to acknowledge the members of my committee, Dr. Wayne Madison and Dr. Mark Smotherman, and the help of Andrew Van Pernis for his work on the CUCsThesis class for L^AT_EX.

I am also very appreciative of my parents, Mike and Phyllis, for their encouragement and support in all of my life's pursuits.

Last, but certainly not least, I owe much of my mental well-being to my wife Corinna, whose love and encouragement have helped me remain sane through the past few years.

Table of Contents

Title Page	i
Abstract	ii
Acknowledgements	iii
Table of Contents	iv
List of Tables	v
List of Figures	vi
List of Algorithms	vii
List of Source Listings	viii
1 Introduction	1
1.1 Background	2
1.1.1 Aggregated Traffic Streams	5
1.2 Related Work	8
2 The Simulator	10
2.1 ON/OFF-Period Distributions	12
2.1.1 Statistical Calculations	13
2.2 Simulator Verification	15
3 Expected Results	16
4 Simulation Results	18
4.1 Total Simulation Time	21
4.2 Number of Sources	27
4.3 Length of Aggregation Interval	33
4.4 Hurst Parameter	38
4.5 Mean ON \neq Mean OFF	45
4.6 Strictly Alternating Sources	47
4.7 Discrete Time	51
4.8 Poisson ON/OFF-Sources	53
4.9 Application of Units	54
5 Conclusions	55
6 Future Work	57
A Complete Simulation Data	58
B Source Code Listing	76
Bibliography	82

List of Tables

4.1	<i>Simulation 1 — Autocorrelation Data of 32 Trials</i>	20
4.2	<i>Simulation 1 — Variance Data of 32 Trials</i>	20
4.3	<i>Simulation ST-1 — Autocorrelation Data of 32 Trials</i>	24
4.4	<i>Simulation ST-2 — Autocorrelation Data of 32 Trials</i>	24
4.5	<i>Simulations ST-1 ST-2 and ST-3 — Variance Data of 32 Trials</i>	24
4.6	<i>Simulation ST-3 — Autocorrelation Data of 5 Trials</i>	25
4.7	<i>Simulation NS-1 — Autocorrelation Data of 32 Trials</i>	28
4.8	<i>Simulation NS-2 — Autocorrelation Data of 32 Trials</i>	29
4.9	<i>Simulation NS-3 — Autocorrelation Data of 32 Trials</i>	29
4.10	<i>Simulations NS-1, NS-2 and NS-3 — Variance Data of 32 Trials</i>	29
4.11	<i>Simulation NS-4 — Autocorrelation Data of 32 Trials</i>	31
4.12	<i>Simulation AT-1 — Autocorrelation Data of 32 Trials</i>	34
4.13	<i>Simulation AT-2 — Autocorrelation Data of 32 Trials</i>	35
4.14	<i>Simulation AT-3 — Autocorrelation Data of 32 Trials</i>	35
4.15	<i>Simulation AT-8 — Autocorrelation Data of 32 Trials</i>	35
4.16	<i>Simulation AT-9 — Autocorrelation Data of 32 Trials</i>	37
4.17	<i>Simulation H-1 — Autocorrelation Data of 32 Trials</i>	39
4.18	<i>Simulation H-2 — Autocorrelation Data of 32 Trials</i>	40
4.19	<i>Simulation H-3 — Autocorrelation Data of 32 Trials</i>	41
4.20	<i>Simulation H-4 — Autocorrelation Data of 32 Trials</i>	42
4.21	<i>Simulation H-5 — Autocorrelation Data of 32 Trials</i>	43
4.22	<i>Simulation H-6 — Autocorrelation Data of 32 Trials</i>	44
4.23	<i>Simulations D-1, D-2, and D-3 — Average Variance of 32 Trials</i>	52
A.1	<i>Key to Simulation Parameter Labels</i>	58
A.2	<i>Simulation Parameters</i>	58
A.3	<i>Key to Variance Data Labels</i>	60
A.4	<i>Complete Variance Data</i>	60
A.5	<i>Key to Autocorrelation Data Labels</i>	62
A.6	<i>Complete Autocorrelation Data</i>	62

List of Figures

1.1	<i>Autocorrelation of DEC-PKT-4 Traced Ethernet Traffic.</i>	4
1.2	<i>The Packet Arrival Count of ON/OFF-Sources over an Interval of Length B</i>	7
4.1	<i>Simulation 1 — Average Autocorrelation of 32 Trials</i>	19
4.2	<i>Simulation 1 — Autocorrelations of 32 Individual Trials</i>	19
4.3	<i>Simulation ST-1 — Average Autocorrelation of 32 Trials</i>	22
4.4	<i>Simulation ST-2 — Average Autocorrelation of 32 Trials</i>	22
4.5	<i>Simulation ST-1 — Autocorrelations for 4 of 32 Individual Trials</i>	23
4.6	<i>Simulation ST-2 — Autocorrelations for 4 of 32 Individual Trials</i>	23
4.7	<i>Simulation ST-3 — Average Autocorrelation of 5 Trials</i>	26
4.8	<i>Simulation ST-3 — Autocorrelations of 5 Individual Trials</i>	26
4.9	<i>Simulations 1, NS-1, NS-2 and NS-3 — Average Autocorrelation of 32 Trials</i>	28
4.10	<i>Simulation NS-4 — Average Autocorrelation of 32 Trials</i>	30
4.11	<i>Simulation NS-4 — Autocorrelations for 8 of 32 Individual Trials</i>	30
4.12	<i>Simulations ST-1, NS-5, NS-6, NS-7, and NS-8 — Average Autocorrelation of 32 Trials</i>	32
4.13	<i>Simulations ST-2, NS-9 and NS-10 — Average Autocorrelation of 32 Trials</i>	32
4.14	<i>Simulations 1, AT-1, AT-2, and AT-3 — Average Autocorrelation of 32 Trials</i>	34
4.15	<i>Simulations AT-4, AT-5, AT-6, and AT-7 — Average Autocorrelation of 32 Trials</i>	36
4.16	<i>Simulations 1 and AT-8 — Average Autocorrelation of 32 Trials</i>	36
4.17	<i>Simulations AT-9 and ST-1 — Average Autocorrelation of 32 Trials</i>	37
4.18	<i>Simulation H-1 — Average Autocorrelation of 32 Trials</i>	39
4.19	<i>Simulation H-2 — Average Autocorrelation of 32 Trials</i>	40
4.20	<i>Simulation H-3 — Average Autocorrelation of 32 Trials</i>	41
4.21	<i>Simulation H-4 — Average Autocorrelation of 32 Trials</i>	42
4.22	<i>Simulation H-5 — Average Autocorrelation of 32 Trials</i>	43
4.23	<i>Simulation H-6 — Average Autocorrelation of 32 Trials</i>	44
4.24	<i>Simulations MO-{1, 2, 3, 4, 5, 6} — Average Autocorrelation of 32 Trials</i>	46
4.25	<i>Simulations MO-{7, 8, 9, 10, 11, 12} — Average Autocorrelation of 32 Trials</i>	46
4.26	<i>Simulations STR-1, STR-2, STR-3, and STR-4 — Average Autocorrelation of 32 Trials</i>	48
4.27	<i>Simulations STR-5, STR-6, STR-7, and STR-8 — Average Autocorrelation of 32 Trials</i>	48
4.28	<i>Simulations STR-9 and H-2 — Average Autocorrelation of 32 Trials</i>	49
4.29	<i>Simulations STR-10 and H-3 — Average Autocorrelation of 32 Trials</i>	49
4.30	<i>Simulations STR-11 and H-5 — Average Autocorrelation of 32 Trials</i>	50
4.31	<i>Simulations STR-12 and AT-7 — Average Autocorrelation of 32 Trials</i>	50
4.32	<i>Simulations D-1, D-2, and D-3 — Average Autocorrelation of 32 Trials</i>	52
4.33	<i>Simulation P-1 — Autocorrelation of 1 Trial</i>	53

List of Algorithms

2.1	<i>The Simulator Algorithm</i>	11
-----	--------------------------------	----

List of Source Listings

B.1 *Simulator Source* 76

Chapter 1

Introduction

Simulation and real-world testing of networking hardware and software such as routers, hubs, and communication protocols is essential for achieving competitive performance. Realistic traffic models are a fundamental requirement for understanding queuing behavior, buffer sizing, congestion management, admission control, and other protocol and hardware design issues.

Tests and simulations are often driven by either simulated traffic or a traffic trace taken from a working network environment. A simulation or test driven by a traffic trace can provide very accurate results when investigating a network with properties similar to the traced network. Unfortunately, large collections of representative traffic traces are difficult to obtain and manage. Given this limitation, the ability to easily synthesize realistic network traffic is a desirable goal. However, many of the currently used methods are either very complex, fail to capture intrinsic properties of real traffic, or both.

One important property that is sometimes overlooked in traffic models is the concept of *long-range dependence* or *self-similarity*, terms introduced by Mandelbrot [1969] in the context of economic patterns. It is shown in [Leland et al. 1994] that Ethernet Local-Area Network (LAN) traffic behaves in a manner similar to that exhibited by a sample from a self-similar, stochastic process. That is, dependencies between traffic utilization levels exist at many different time scales. This means that a period of network utilization above (or below) the mean level is likely to remain above (or below) the mean for an extended period of time. Dependencies remain between utilization levels that are measured milliseconds, minutes or even hours apart.

Taqqu et al. [1997] describe a simple method for generating a network traffic stream, and prove that with an infinite number of sources and an infinite amount of time, the generated stream is self-similar (implying that it realistically models real-world traffic). Unfortunately, infinite limits are

inherently *unrealistic*, therefore, our work focuses on determining what *finite* limits give reasonable results.

The effects of self-similar traffic on networking environments are not well understood, despite numerous papers on the subject. Park et al. [1997] and Erramilli et al. [1996] conclude that self-similar traffic can have a serious adverse impact on network performance, and that ignoring the effects of self-similarity typically results in overly optimistic performance predictions. On the other hand, Cao et al. [1999] and Christiansen et al. [2000] provide simulation evidence that self-similar traffic tends to behave much like traffic that does not have significant dependencies, especially at high network utilizations.

Though a discussion on the effects of self-similar traffic on network performance is beyond the scope of this thesis, the results of this work help provide an intuition into simulating traffic with various degrees of dependence at the packet level. This could aid researchers in designing simulations which use self-similar traffic models, allowing further study of the effects of these models on network performance.

1.1 Background

In [Leland et al. 1994], traces of Ethernet LAN traffic were analyzed and found to be self-similar in nature. That is, a graph of packet arrivals per unit time from the examined Ethernet traffic tends to look “similar” when viewed over a wide range of time scales — from milliseconds to hours. In particular, there is no typical length of a traffic *burst*, and the variance of the packet arrivals per unit time appears large.

Statistically, the property that characterizes self-similarity, or more precisely, long-range dependence, is a slowly decaying autocorrelation function. The autocorrelation, $r(k)$, of a stochastic

process, X_t , with mean μ and variance σ^2 is defined as

$$r(k) = E[(X_t - \mu)(X_{t+k} - \mu)] / \sigma^2. \quad (1.1)$$

The autocorrelation is a measure of the correlation between elements of X_t that are a distance of k apart. The values of $r(k)$ are normalized such that if all elements a distance of k apart are nearly identical, $r(k) \rightarrow 1$. If the elements of X_t are independent and identically distributed (i.i.d.), $r(k) = 0$ for $k \geq 1$. [Geist and Westall 2000]

For a stochastic process to be long-range dependent, it must have a non-summable autocorrelation, meaning $\sum_k r(k) = \infty$. Exactly self-similar processes have autocorrelation functions that satisfy

$$r(k) = \frac{1}{2} [(k+1)^{2H} - 2k^{2H} + (k-1)^{2H}] \quad (1.2)$$

where $H : 0.5 < H < 1.0$ is the *Hurst parameter*, or degree of self-similarity. Exactly self-similar processes with H in this range are long-range dependent.

It is important to note that self-similarity or long-range dependence cannot be verified for a finite sample. Following the terminology used in [Leland et al. 1994], it can be said (somewhat loosely) that a finite sample is self-similar *in nature* if it is statistically consistent with a sample of a stochastic process that is self-similar.

Thus, after statistically analyzing several traces of Ethernet LAN traffic, Leland et al. [1994] claim that the traces are self-similar in nature, with an estimated Hurst parameter $H \approx 0.80$. See Figure 1.1 for a plot of the autocorrelation of a publicly available Ethernet trace used in this study. This plot shows the autocorrelation calculated for the packet arrival counts of the 14,400 adjacent, non-overlapping 0.25 second blocks in the one hour trace. The autocorrelation appear to be slowly-decaying, and closely follows the exactly self-similar trend of Equation 1.2. Traces of traffic at NSFNET core switches were analyzed by Klivansky et al. [1994] and determined to behave in a long-range dependent manner at the packet level as well, with a mean Hurst parameter $H = 0.74$.

Additional studies show that traces of Wide-Area Network (WAN) traffic [Paxson and Floyd 1995], Variable Bit Rate (VBR) video traffic [Beran et al. 1995] and World-Wide Web (WWW) traffic [Crovella and Bestavros 1997] are also self-similar in nature.




Figure 1.1: *Autocorrelation of DEC-PKT-4 Traced Ethernet Traffic.*
Available at <http://ita.ee.lbl.gov/html/contrib/DEC-PKT.html>

There are many techniques for synthesizing processes that exhibit self-similarity, and are therefore suitable for generating network traffic. Among the most common are numerical methods for generating *fractional Gaussian noise* (fGn) [Paxson 1995; Che and Li 1997; Geist and Westall 2000]. fGn is the increment process of *fractional Brownian motion* (FBM), which is also called the *Random Walk Process*. fGn has a normal distribution with mean $\mu = 0$, and variance $\sigma^2 \propto B^{2H}$, where B is the width of the FBM increment. fGn is exactly self-similar with an autocorrelation given by Equation 1.2. A stream of packet arrival counts per unit time can be synthesized from fGn by specifying a mean utilization level and using fGn as a deviation from that mean.

While the use of these techniques, and other methods described in [Popescu 1999], can accurately

simulate the self-similar nature of network traffic, these approaches are undesirable according to Willinger et al. [1997] and Nikolaidis et al. [1997] for a number of reasons. First, they do not provide any explanation for the cause of long-range dependence observed in network traffic. Second, these techniques typically exhibit at least $O(n \log n)$ run-time complexity in terms of the length of time simulated. Third, they often are not parallelizable with near linear speedup. Finally, some techniques require a large number of parameters to achieve the desired statistical characteristics.

In an attempt to address these problems, Leland et al. [1994] propose modeling the traffic streams emanating from individual sources. Each source is either ON or OFF at any given time. Given a realistic model of the traffic from a single source, one would expect that by aggregating the ON/OFF streams of many sources together, the resulting traffic would have characteristics similar to that of traced network traffic. The key is determining the length of time that each source spends in an ON or OFF state. Simulations that sample these times from a Poisson process (which has an exponentially decaying probability density function) fail to capture the self-similar nature of real traffic. Paxson and Floyd [1995] show that while the Poisson model holds for certain types of traffic (e.g. session and connection arrivals, and TELNET traffic), other types of traffic do not fit well with a Poisson model. By taking the ON/OFF times from a heavy-tailed distribution with an infinite variance (such as the Pareto distribution), it is shown in [Leland et al. 1994] (and proven in [Taqqu et al. 1997]) that it is theoretically possible to generate self-similar traffic by aggregating multiple sources. To understand the concept of this aggregated traffic stream, it is useful to follow the description as laid out in [Willinger et al. 1997].

1.1.1 Aggregated Traffic Streams

Consider a variation of the discrete *reward renewal process* used by Mandelbrot [1969] to generate fGn, also known as the *packet train model* in the context of network traffic [Jain and Routhier 1986]. Since an individual source is either ON or OFF at a given time, let $W(t)$, $t \geq 0$ represent the state of that source at time t , where $W(t) = 1$ and $W(t) = 0$ indicate the existence or absence, respectively,

of a packet at time t . During an ON-period ($t_1 \leq t \leq t_2$), when the source is generating traffic, $W(t)$ represents a “reward” of 1, and during an OFF-period when the source is not generating traffic, $W(t)$ represents a “reward” of 0.¹ The lengths of both the ON- and OFF-periods are independent and identically distributed (i.i.d.), and the ON- and OFF-periods are independent of each other. We consider both the case where the length of the ON/OFF-periods are from identical distributions, and the case where the distributions have different means. We also consider both *strictly alternating* sources, where an ON-period is immediately followed by an OFF-period (and vice-versa), and *idealized* sources, where an ON- or OFF-period is just as likely to be followed by another ON- or OFF-period.

Now consider M sources, and let $W^{(m)}(t)$ be the state of source m at time t , where $m = 1, \dots, M$. The total packet arrival count at time t , $W_M^*(t)$, is then:

$$W_M^*(t) = \sum_{m=1}^M W^{(m)}(t) \quad (1.3)$$

By aggregating over a time block of length B , we arrive at the cumulative packet arrival count over the range $[0, Bt]$:

$$W_M^*(Bt) = \int_0^{Bt} \left(\sum_{m=1}^M W^{(m)}(u) \right) du \quad (1.4)$$

The increment of the cumulative arrival count, $W'(t) = W_M^*(B(t+1)) - W_M^*(Bt)$, represents the aggregated packet arrival count over the interval $[Bt, B(t+1)]$, which has interval length B . The aggregation of three ON/OFF-sources is demonstrated in Figure 1.2.

For large M and B (that is, as $M \rightarrow \infty$ and $B \rightarrow \infty$), the behavior of Equation 1.4 depends upon the distributions of the lengths of the ON/OFF-periods. If (at least one of) the ON/OFF-processes is heavy-tailed, $W_M^*(Bt)$ will converge to FBM, and hence $W'(t)$ will converge to fGn.

¹To make the model more intuitive, it is possible to assume that the granularity of $W(t)$ is smaller than the time necessary to send a single packet on a physical Ethernet, for example. The state of $W(t)$ will be ON during the time required to transmit a packet, and OFF during the inter-packet time. This restriction is not necessary, however. The model holds when an ON-period represents traffic sent at a constant rate over a larger time scale.




Figure 1.2: *The Packet Arrival Count of ON/OFF-Sources Aggregated over an Interval of Length B*

Heavy-tailed distributions have infinite variance, or more intuitively, they allow very large lengths of ON/OFF-periods with non-negligible probability. ON/OFF-periods modeled by Poisson processes do *not* have infinite variance, and will *not* cause $W'(t)$ to converge to fGn with large M and B . Instead, they will cause $W'(t)$ to converge to white noise, which lacks any significant correlation.

Assume that the complementary distributions of the ON- and OFF-periods, respectively, can be denoted by the following (heavy-tailed) functions:

$$F_{1c}(x) \sim l_1 x^{-\alpha_1} L_1(x) \quad \text{with} \quad 1 < \alpha_1 < 2 \quad (1.5)$$

$$F_{2c}(x) \sim l_2 x^{-\alpha_2} L_2(x) \quad \text{with} \quad 1 < \alpha_2 < 2$$

Then, by letting

$$a_j = l_j \frac{\Gamma(2 - \alpha_j)}{\alpha_j - 1} \quad (1.6)$$

$$b = \lim_{t \rightarrow \infty} t^{\alpha_2 - \alpha_1} \frac{L_1(t)}{L_2(t)} \quad (1.7)$$

the variance, σ_{lim}^2 , of the aggregated packet count, Equation 1.4, can be calculated as follows:

$$\sigma_{lim}^2 = \frac{2(\mu_2^2 a_1 b + \mu_1^2 a_2)}{(\mu_1 + \mu_2)^3 \Gamma(4 - \alpha_{min})} \quad (1.8)$$

The autocorrelation (Equation 1.1) of the aggregated packet count will be exactly self-similar, of the form in Equation 1.2, where $\alpha = 3 - 2H$.

The benefits of this method are that it is conceptually and mathematically simple (even if the proof of self-similarity is not), a simulation can be generated by specifying only a few parameters, it is inherently parallel, the runtime increases linearly with the length of the simulation time (given a fixed number of sources), and finally, it provides an intuition into the cause of the self-similar behavior of traffic.

Critics question the ability of this simple model to accurately describe all the nuances of real traffic, since many factors that affect traffic, such as routing queues, congestion and congestion avoidance, packet losses, Ethernet collisions, user interaction, etc., are not each taken into consideration. Each of these factors can influence the burstiness and dependence of traffic, and each at a different time scale, from short-term *spikes* to long-term *swells*. The argument can be made however, that the fractal nature of a self-similar, fGn model captures all of these dependencies over a wide range of time scales, eliminating the need to describe them individually.

1.2 Related Work

Simulating traffic using an infinite variance ON/OFF source model as outlined in Section 1.1 was performed in [Taqqu et al. 1997]. However, the results of only a few tests were mentioned, all were run using a large number of sources (either 500 or 16,000), and no statistical results were given. They claimed that with a 16,384 node MasPar MP-1216, network traffic that is visually similar to traced traffic could be quickly simulated. Indeed, visual similarity was demonstrated between the simulated and traced traffic, and large visual differences were demonstrated between the simulated traffic and traffic generated with the more traditional Poisson ON/OFF model.

Christiansen et al. [2000] analyzed the performance of Random Early Detection (RED) in routers by using between 700 and 5,075 simulated HTTP users. The *user think times* (OFF-periods) and

response sizes (ON-periods) were both sampled from heavy-tailed Pareto distributions. Though this study demonstrates the performance effects of traffic generated with the infinite-variance ON/OFF-source model, it provides no analysis of the actual generated traffic.

A parallel, message-passing algorithm is presented in [Nikolaidis et al. 1997] for synthesizing ATM traffic that is self-similar in nature. They demonstrate that it is possible to use the ON/OFF-source model to generate ATM traffic in real-time for a 155 Mbps link. While simulating 500 sources, they achieved near-linear speedup as the number of workstations increased (up to 8), at least for a link utilization of 20%. They do not give any statistical results for the synthesized traffic, nor do they provide any insight into the number of sources that might be required to generate traffic that is self-similar in nature. But because the parameters used in their tests were comparable to those used in the simulations presented in [Taqqu et al. 1997], we can conclude that the traffic generated by Nikolaidis et al. [1997] would also appear pictorially similar to traced Ethernet traffic. Thus, their work demonstrates the feasibility of synthesizing ATM traffic that appears self-similar in nature, while using a relatively small network of machines (e.g. 8).

These studies, however, fail to identify the ranges of values for the number of sources, M , and the size of the observation interval, B , that result in simulated traffic that *reasonably* approximates the self-similar nature of fractional Gaussian noise (fGn).² The range and sensitivity of these limits is the focus of our work. We are also interested in determining if significant differences are seen when using strictly alternating or idealized ON/OFF-sources, and when setting the mean length of the ON- and OFF-processes in such a way that the mean number of sources that are ON at any given time (i.e. total link utilization) is ≤ 1.0 .

²Note that we have not yet defined *reasonable*.

Chapter 2

The Simulator

The simulator algorithm, which is derived from Equation 1.4, generates a stream of aggregated packet arrival counts given the following input parameters:

H — the target Hurst parameter (degree of self-similarity)

M — the number of ON/OFF-sources

μ_{on} — the mean length of an ON-period

μ_{off} — the mean length of an OFF-period

B — the time interval over which the ON/OFF traffic streams are aggregated to produce one packet arrival count

T — the total length of time to simulate

We assume a fixed packet size for the simulation, thus packet or byte arrival counts differ only by a constant factor and are considered equivalent for our purposes. The number of packet arrivals per interval of length B is then the sum of the time each source spends in an ON state during that interval. For every interval, each source is “run” through its ON/OFF sequence, and the ON times are recorded, until the end of the most recent ON- or OFF-period extends beyond the end of the current interval. The simulator algorithm is presented in Algorithm 2.1.

To begin data collection only after the simulation has reached a steady state, the start time of each source, $t_0[i]$, is uniformly distributed such that $-10M(\mu_{on} + \mu_{off}) \leq t_0[i] < 0$. The source is run until the end of its most recent ON- or OFF-period, $time[i]$, becomes positive. The recording of ON-times does not begin until time zero.

Though the mathematical model presented in Section 1.1 assumes a discrete time model, the simulator actually implements a continuous time model, at least as much as the precision of a 64-bit `double` allows. That is, a real number is sampled from the appropriate random distribution for

Algorithm 2.1 *The Simulator Algorithm*

```

1: {Initialize sources and run them until time zero}
2: for  $i \leftarrow 1$  to  $M$  do
3:    $time[i] \leftarrow$  initial start time  $\leq 0$ 
4:    $state[i] \leftarrow$  OFF
5:
6:   while  $time[i] \leq 0$  do
7:     if  $state[i]$  is OFF then
8:        $time[i] \leftarrow time[i] +$  length of next ON-period
9:        $state[i] \leftarrow$  ON
10:    else
11:       $time[i] \leftarrow time[i] +$  length of next OFF-period
12:       $state[i] \leftarrow$  OFF
13:    end if
14:  end while
15: end for
16:
17: {For each interval, run each source, recording ON-times, until the end of the interval}
18: for  $t \leftarrow 1$  to  $T/B$  do
19:    $interval\_start \leftarrow (t - 1) \times T$ 
20:    $interval\_end \leftarrow t \times T$ 
21:    $on\_time[t] \leftarrow 0$ 
22:
23:   for  $i \leftarrow 1$  to  $M$  do
24:     if  $state[i]$  is ON then
25:       {Record ON time since beginning of interval}
26:        $on\_time[t] \leftarrow on\_time[t] + \min(B, time[i] - interval\_start)$ 
27:     end if
28:
29:     while  $time[i] \leq interval\_end$  do
30:       if  $state[i]$  is OFF then
31:          $p \leftarrow$  length of next ON-period
32:          $time[i] \leftarrow time[i] + p$ 
33:         {Record ON time, up to end of interval}
34:          $on\_time[t] \leftarrow on\_time[t] + \min(p, interval\_end - time[i])$ 
35:          $state[i] \leftarrow$  ON
36:       else
37:          $p \leftarrow$  length of next OFF-period
38:          $time[i] \leftarrow time[i] + p$ 
39:          $state[i] \leftarrow$  OFF
40:       end if
41:     end while
42:   end for
43: end for

```

the length of the ON- or OFF-period of a source, and that number is used to advance the “current time” of that source. It is not necessary that the source begin and end an ON- or OFF-period on a particular discrete boundary. This means that the packet arrival count for a given interval of width B can be a non-integer value.

In addition to the parameters mentioned above, the following options are also supported to change the behavior of the algorithm:

strict/notstrict — causes the sources to strictly alternate between an ON-period and an OFF-period, or causes the likelihood of the next period being either ON or OFF to occur with equal probability, respectively

discrete — allows the specification of a discretization value (as opposed to the precision of a 64-bit **double**)

Algorithm 2.1 is strictly alternating. To accomplish idealized sources, an ON or OFF state is chosen (with equal probability) between lines 6 and 7, and between lines 29 and 30. In order to make the simulator *discrete*, values are truncated to the specified discretization level at lines 3, 8, 11, 31, and 37. It is assumed that the discretization level is a factor of B , and that B is a factor of T .

The run-time complexity of the simulator is linear over the number of values that must be generated from a Pareto distribution. In the average case, the simulator has a run-time complexity of $\Theta(\frac{M \times T}{\mu_{on} + \mu_{off}})$.

2.1 ON/OFF-Period Distributions

The lengths of the ON- and OFF-periods are both sampled from a Pareto distribution, which is specified by the probability density function

$$p(x) = \alpha \beta^\alpha x^{-(\alpha+1)}, \quad x \geq \beta \quad (2.1)$$

The cumulative distribution function of a Pareto is

$$P(x) = 1 - \left(\frac{\beta}{x}\right)^\alpha, \quad \alpha, \beta \geq 0, x \geq \beta \quad (2.2)$$

and the mean is given by

$$\mu = \frac{\beta\alpha}{\alpha - 1}. \quad (2.3)$$

For our purposes, the inverse cumulative distribution function,

$$P_{-1}(r) = \beta(1 - r)^{\frac{-1}{\alpha}} \quad (2.4)$$

is used to map a random number r , uniformly distributed over $[0, 1)$, onto the Pareto distribution.

Following Taqqu et al. [1997] and Equation 2.3, we can set α and β as

$$\alpha = 3 - 2H \quad (2.5)$$

$$\beta = \frac{\mu(\alpha - 1)}{\alpha} \quad (2.6)$$

allowing us to specify the distribution in terms of only H and μ , the target Hurst parameter, and the mean length of a source's ON or OFF period.

For a given simulation using this model, the lengths of the ON- and OFF-periods are obtained using the same value of H , implying that α_{on} will always equal α_{off} . For the case of identical ON/OFF-processes, $\mu_{on} = \mu_{off}$ as well.

2.1.1 Statistical Calculations

Once the simulator has generated the aggregate packet arrival process, $\hat{W}'(t)$, its autocorrelation and variance are then computed.

It is stated in Theorem 1 of [Taqqu et al. 1997] that for large M and B , the cumulative packet

arrival count, $W_M^*(Bt)$, behaves statistically like

$$BM \frac{\mu_{on}}{\mu_{on} + \mu_{off}} t + B^H \sqrt{M} \sigma_{lim} B_H(t) \quad (2.7)$$

where $E[W_M^*(Bt)] = BM \frac{\mu_{on}}{\mu_{on} + \mu_{off}} t$, the term $B_H(t)$ is fractional Brownian motion (FBM), which provides deviations from the expected value, scaled by a factor of $T^H \sqrt{M} \sigma_{lim}$. More precisely,

$$\mathcal{L} \lim_{B \rightarrow \infty} \mathcal{L} \lim_{M \rightarrow \infty} \frac{\left(W_M^*(Bt) - BM \frac{\mu_{on}}{\mu_{on} + \mu_{off}} t \right)}{B^H \sqrt{M}} = \sigma_{lim} B_H(t) \quad (2.8)$$

where $\mathcal{L} \lim$ means convergence in the sense of the finite dimensional distributions [Taqqu et al. 1997]. Taking the derivative with respect to t yields

$$\mathcal{L} \lim_{B \rightarrow \infty} \mathcal{L} \lim_{M \rightarrow \infty} \frac{\left(W'(t) - BM \frac{\mu_{on}}{\mu_{on} + \mu_{off}} \right)}{B^H \sqrt{M}} = \sigma_{lim} G_H(t) \quad (2.9)$$

where $G_H(t)$ is fractional Gaussian noise.

In similar fashion to Equation 2.9, the normalized variance, σ_{norm}^2 , of a sample aggregate packet arrival count, $\hat{W}'(t)$, can be calculated as follows:

$$\sigma_{norm}^2 = \sum_{t=1}^N \frac{(\hat{W}'(t) - \mu_{exp})^2}{B^{2H} M N} \quad (2.10)$$

where $\mu_{exp} = BM \frac{\mu_{on}}{\mu_{on} + \mu_{off}}$ and $N = T/B$.

The sample autocorrelation of $\hat{W}'(t)$ is

$$r(k) = \frac{1}{N-k} \sum_{t=1}^{N-k} \frac{(\hat{W}'(t) - \mu_{exp})(\hat{W}'(t+k) - \mu_{exp})}{\sigma_{calc}^2} \quad (2.11)$$

where σ_{calc}^2 the calculated, non-normalized sample variance of $\hat{W}'(t)$.

This autocorrelation, $r(k)$, along with the normalized variance, σ_{norm}^2 , are examined to de-

termine, with respect to various parameters, the effectiveness of the ON/OFF-source model at approximating exactly self-similar traffic.

2.2 Simulator Verification

No formal analysis has been performed to verify that the simulator correctly implements the mathematical model presented in Section 1.1. However, results of a small set of simulations have been found to be consistent with results obtained from an independently generated simulator, and the results match our expectations in most cases. In addition, the mean and variance of samples generated from Pareto distributions have been shown to agree with the mean and variance of the distributions. Though we can never eliminate the possibility of implementation errors, these techniques, combined with careful hand-checking of the simulator greatly increase the probability that all implementation errors have been corrected.

Chapter 3

Expected Results

Based on the work of Taqqu et al. [1997], it is known to be possible to generate network traffic that is self-similar in nature using the aggregated ON/OFF-stream model. Their simulations used the parameters: Hurst parameter $H = 0.9$, number of sources $M = 500$, identical, strictly alternating ON/OFF-sources, and a total simulation time of 27 hours at a time-scale of 10 milliseconds.

We would expect that as the limits of the simulated packet arrival count, $\hat{W}'(t)$, increase (that is, the number of sources, M , and the scale of the aggregation interval, B), the resulting autocorrelations would converge to the target: exactly self-similar as in Equation 1.2. We would also expect that as M and B decrease, the resulting autocorrelations would diverge from the target.

For any reasonable M , B , and H , we anticipate that as either the length of the simulation time, T , is increased, or the autocorrelations of many smaller trials runs are averaged, the resulting autocorrelation values will converge to points that lie on a smooth curve (though not necessarily the target).

As the limiting parameters of $\hat{W}'(t)$ increase, the resulting normalized variance, σ_{norm}^2 , should converge to the value given in Equation 1.8. For i.i.d. ON/OFF-processes, the expected, normalized variance can be further simplified to

$$\sigma_{lim}^2 = \frac{\beta^\alpha}{2\mu(\alpha - 1)(2 - \alpha)(3 - \alpha)} \quad (3.1)$$

where $0.5 < H < 1.0$, $\alpha = 3 - 2H$ and $\beta = \frac{\mu(\alpha-1)}{\alpha}$.

For heavy-tailed ON/OFF-processes with the same α but different means, μ_{on} and μ_{off} , σ_{norm}^2

is expected to converge to

$$\sigma_{lim}^2 = \frac{2(\mu_{off}^2 \beta_{on}^\alpha + \mu_{on}^2 \beta_{off}^\alpha)}{(\mu_{on} + \mu_{off})^3 (\alpha - 1)(2 - \alpha)(3 - \alpha)} \quad (3.2)$$

as M and B become large.

In terms of the limits of the aggregation time, B , we would expect the ratio of B to the mean ON/OFF-lengths and total simulation time to be more significant than the actual value of B . This implies that a simulation with $\mu_{on} = \mu_{off} = 1.0$, $B = 1000$ and a total simulation time of $T = 10^9$ should yield identical results to a simulation with $\mu_{on} = \mu_{off} = 0.1$, $B = 100$, and $T = 10^8$. This allows the units to be interpreted as seconds, milliseconds, or at any other time scale, with equivalent results.

Our objective is to determine if values for M and B exist such that real-world synthesis of traffic that is self-similar in nature is feasible with a small network of machines in a reasonable amount of time.

Chapter 4

Simulation Results

To demonstrate the feasibility of simulating traffic that is self-similar in nature, while using a significantly smaller number of sources than presented in [Taqqu et al. 1997], we use a simulation with the parameters $H = 0.75$, $M = 32$ sources, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$, and idealized ON/OFF-sources as a baseline. These parameters were identified experimentally. This represents about 28.4 hours worth of traffic where the mean ON-period for each source is one millisecond, and the autocorrelation is calculated using the packet arrival counts for 1024 millisecond intervals. Figure 4.1 shows the average autocorrelation of the 32 trials (shown individually in Figure 4.2) that were run with these parameters. The average is visually very close to the exactly self-similar target for $H = 0.75$. Each of the 32 trials took one hour to simulate on a machine with a 742 MHz Pentium II Xeon processor.

Table 4.1 shows the visual similarity in a more quantitative way. The *Lag* is the distance, k , between packet arrival counts whose correlation is calculated, $\hat{\mu}_{r(k)}$ is the mean sample autocorrelation value and $\Delta\mu$ is the difference between $\hat{\mu}_{r(k)}$ and the target value. The column labeled $\hat{\sigma}_{r(k)}$ is the standard deviation of $r(k)$ for the 32 trials, and $\pm 90\%$ *C.I.* represents the half-width of the two-sided 90% confidence interval. In all but one case for Simulation 1, the target value is contained within the 90% confidence interval.

The variance of the packet arrival count for Simulation 1 is shown in Table 4.2, where σ_{lim}^2 is the expected variance from Equation 3.1, $\hat{\sigma}_{norm}^2$ is the mean, normalized sample variance, and $\Delta\sigma^2$ is the difference between the two. The column labeled $\hat{\sigma}_{\sigma^2}$ is the standard deviation of the 32 normalized sample variances. We note that the target variance is not within the 90% confidence interval, though the difference between the target and mean sample variance is quite small.

Quantifying the quality of a simulation's results requires a bit of comment. It is most desirable




Figure 4.1: *Simulation 1 — Average Autocorrelation of 32 Trials*
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$




Figure 4.2: *Simulation 1 — Autocorrelations of 32 Individual Trials*
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$

Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
1	0.4149	-0.00069	0.0125	0.00364	0.4142	Yes
2	0.2668	0.00288	0.0156	0.00455	0.2696	Yes
4	0.1841	0.00412	0.0171	0.00499	0.1882	Yes
8	0.1284	0.00427	0.0176	0.00511	0.1327	Yes
16	0.0890	0.00480	0.0186	0.00542	0.0938	Yes
32	0.0614	0.00489	0.0179	0.00520	0.0663	Yes
64	0.0431	0.00378	0.0188	0.00548	0.0469	Yes
128	0.0304	0.00277	0.0195	0.00568	0.0331	Yes
256	0.0195	0.00396	0.0176	0.00513	0.0234	Yes
512	0.0111	0.00546	0.0175	0.00510	0.0166	No

Table 4.1: *Simulation 1 — Autocorrelation Data of 32 Trials*
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$

σ_{lim}^2	$\hat{\sigma}_{norm}^2$	$\Delta\sigma^2$	$\hat{\sigma}_{\sigma^2}$	$\pm 90\%$ C.I.	Target in C.I.
0.2566	0.2527	0.00393	0.0050	0.00147	No

Table 4.2: *Simulation 1 — Variance Data of 32 Trials*
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$

for the autocorrelation values to lie exactly on the target line. In practice, though, this rarely happens. In the following sections, we will say that a simulation provides “good” or “reasonable” results if the target is contained within the 90% confidence interval for every, or all except one, value of the lag. However, to compare two simulations who’s 90% confidence intervals contain the target in every case, we need to closely examine both the distance of the mean autocorrelations from the target and the standard deviation of the sample autocorrelations. Thus, claiming that the results of a simulation are “good” or “better” than another simulation remains somewhat subjective.

Nonetheless, we believe that Simulation 1 represents a good baseline to which we compare further simulations. It can also be claimed that Simulation 1 provides “good” results in terms of the self-similar nature of the simulated traffic stream. Now we modify various parameters in turn to examine their limits and sensitivity.

4.1 Total Simulation Time

The first parameter we adjust is the length of the simulation time, T . A short length, that still provides good results, is the most desirable, since longer simulations can take a large amount of clock-time to run.¹ It is desirable to have the results converge to the target values within a short T when using this method to synthesize traffic on a real network as well, since there is no way to speed up the actual clock-time compared to the simulation time in this case.

We begin by reducing the simulation time of Simulation 1 by a factor of 10 for Simulation ST-1 (Figure 4.3), and then by another 10 for Simulation ST-2 (Figure 4.4). We immediately notice three things about these simulations. First, the average autocorrelation of Simulation ST-1 is not nearly as close to the target as Simulation 1, and Simulation ST-2 is even worse. Second, higher values of the lag, k , result in a less smooth autocorrelation. This is especially evident in Figures 4.5 and 4.6, where a representative subset of the 32 individual trials of Simulations ST-1 and ST-2, respectively, are shown. We argue that this decreased smoothness is largely due to the low number of data points, $N = T/B$, used in the autocorrelation calculation. Third, Figures 4.5 and 4.6 show how the autocorrelations of the individual trials deviate from the average by an increasing amount as the T decreases.

Tables 4.3 and 4.4 show these observations numerically. The target value is almost never in the 90% confidence interval for the sample autocorrelations of Simulations ST-1 and ST-2. Interestingly, the standard deviation of the autocorrelations for Simulation ST-1 are almost twice that of Simulation 1, and Simulation ST-2 is more than twice that of Simulation ST-1.

The variance of Simulations ST-1 and ST-2 shown in Table 4.5, also deviates from the target value as T decreases, and the standard deviation of these calculated variances increases twofold for Simulation ST-1, and another threefold for Simulation ST-2. The calculated variance, even for Simulation ST-2, is still quite close to the target value, however.

¹*Clock-time* is the amount of time a simulation takes to run as measured by a wall clock or watch, as opposed to *simulation time* which is the amount of time simulated — typically much larger than the clock-time.




Figure 4.3: *Simulation ST-1 — Average Autocorrelation of 32 Trials*
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^4$




Figure 4.4: *Simulation ST-2 — Average Autocorrelation of 32 Trials*
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^3$




Figure 4.5: *Simulation ST-1* — Autocorrelations for 4 of 32 Individual Trials
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^4$




Figure 4.6: *Simulation ST-2* — Autocorrelations for 4 of 32 Individual Trials
 $H = 0.75$, $M = 32$, $B = 1024$, $\mu_{on} = \mu_{off} = 1$, $T = 1024 \times 10^3$

Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
1	0.4060	0.008173	0.0248	0.007204	0.4142	No
2	0.2553	0.014318	0.0269	0.007828	0.2696	No
4	0.1712	0.017025	0.0351	0.010200	0.1882	No
8	0.1060	0.026719	0.0355	0.010319	0.1327	No
16	0.0725	0.021254	0.0355	0.010320	0.0938	No
32	0.0424	0.023892	0.0302	0.008772	0.0663	No
64	0.0210	0.025835	0.0292	0.008483	0.0469	No
128	0.0129	0.020291	0.0244	0.007098	0.0331	No
256	0.0024	0.020989	0.0211	0.006144	0.0234	No
512	0.0028	0.013723	0.0159	0.004621	0.0166	No

Table 4.3: *Simulation ST-1 — Autocorrelation Data of 32 Trials*
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^4$

Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
1	0.3826	0.031644	0.0668	0.019439	0.4142	No
2	0.2186	0.051057	0.0778	0.022615	0.2696	No
4	0.1405	0.047796	0.0832	0.024208	0.1882	No
8	0.0851	0.047612	0.0839	0.024392	0.1327	No
16	0.0537	0.040033	0.0933	0.027126	0.0938	No
32	0.0331	0.033182	0.0874	0.025413	0.0663	No
64	0.0095	0.037400	0.0810	0.023558	0.0469	No
128	0.0147	0.018430	0.0819	0.023807	0.0331	Yes
256	0.0148	0.008640	0.0744	0.021635	0.0234	Yes

Table 4.4: *Simulation ST-2 — Autocorrelation Data of 32 Trials*
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^3$

Sim	σ_{lim}^2	$\hat{\sigma}_{norm}^2$	$\Delta\sigma^2$	$\hat{\sigma}_{\sigma^2}$	$\pm 90\%$ C.I.	Target in C.I.
ST-1	0.2566	0.2492	0.00304	0.0105	0.00738	No
ST-2	0.2566	0.2438	0.00898	0.0309	0.01282	No
ST-3	0.2566	0.2538	0.00279	0.0025	0.00169	No

Table 4.5: *Simulations ST-1 ST-2 and ST-3 — Variance Data of 32 Trials*
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = \{1024 \times 10^4, 1024 \times 10^3, 1024 \times 10^6\}$

Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
1	0.4180	-0.00375	0.0074	0.00548	0.4142	Yes
2	0.2717	-0.00202	0.0086	0.00635	0.2696	Yes
4	0.1893	-0.00109	0.0092	0.00677	0.1882	Yes
8	0.1343	-0.00161	0.0096	0.00703	0.1327	Yes
16	0.0951	-0.00135	0.0096	0.00710	0.0938	Yes
32	0.0675	-0.00124	0.0107	0.00785	0.0663	Yes
64	0.0477	-0.00085	0.0099	0.00726	0.0469	Yes
128	0.0335	-0.00031	0.0105	0.00769	0.0331	Yes
256	0.0228	0.00065	0.0093	0.00688	0.0234	Yes
512	0.0149	0.00166	0.0088	0.00650	0.0166	Yes

Table 4.6: *Simulation ST-3 — Autocorrelation Data of 5 Trials*
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^6$

One of the individual trials in Simulation ST-2 has a surprisingly high correlation, as evident from Figure 4.6. This is likely due to the increased probability that the length of one or more ON-periods chosen from the Pareto distribution would span the entire simulation time.

Finally, Simulation ST-3 was run with a longer simulation time, $T = 1024 \times 10^6$. Each trial took nearly 10 hours to run on a 742 MHz Pentium II Xeon, but as evident from Figures 4.7 and 4.8, and Table 4.6, the average autocorrelation is very close to the target and the standard deviation of the trials is low. The 90% confidence interval is included in Table 4.6 for consistency, though it should be noted that the confidence interval is not an accurate estimate for only 5 trials. However, because both the standard deviation and the distance between autocorrelation values and the target are significantly smaller than the other simulations, we claim that Simulation ST-3 generates the “best” results of any simulation up to this point. The variance is also closer to the target than any other simulation so far (Table 4.5).

Of this small sample of simulation times, $T = 1024 \times 10^5$ seems to give autocorrelations and variances that converge to the target values within a reasonable amount of clock-time. The shape of the autocorrelation for Simulation ST-1 where $T = 1024 \times 10^4$ is also similar to the target, but the values are an order of magnitude farther from the target than Simulation 1.




Figure 4.7: *Simulation ST-3 — Average Autocorrelation of 5 Trials*
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^6$




Figure 4.8: *Simulation ST-3 — Autocorrelations of 5 Individual Trials*
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^6$

4.2 Number of Sources

We have seen that 32 sources is adequate for generating self-similar traffic, which is much more tractable than the 500 (or 16,000) sources used in [Taqqu et al. 1997]. The next step is to investigate the effects of more and fewer sources on the autocorrelation and variance. We continue to use the parameters used in Simulation 1 for $H = 0.75$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$, and idealized ON/OFF-sources.

For Simulations NS-1, NS-2, and NS-3, M is set to 8, 4, and 2 sources, respectively. As shown in Figure 4.9, the number of sources makes only a slight difference. In fact, Table 4.7 might indicate an improvement for 8 sources over 32, since the 90% confidence interval for Simulation NS-1 contains the target values in every case. On close examination, the standard deviation of the autocorrelation values in Simulation NS-1 increased significantly over Simulation 1, creating a wider confidence interval, so an overall improvement cannot be inferred. Simulations NS-1, NS-2, and NS-3 *do* show that given a large enough B and T , traffic with autocorrelations closely matching the target can be generated with as few as 8 sources. Even with as 4 or 2 sources, traffic can be generated with autocorrelations that are surprisingly close to the target. In all three simulations, the average calculated variance is still close to the target variance, though it begins to decrease slightly with fewer sources.

An interesting effect is observed when we set M to 1 source as in Simulation NS-4. As shown in Figure 4.10 and Table 4.11, the autocorrelation is actually slightly *larger* than the target, and the target is contained in the 90% confidence interval for all values of the lag, k . However, as we see in Table 4.11, and graphically in Figure 4.11, the standard deviation of the individual trials in Simulation NS-5 is very large — much larger than the standard deviation for most other simulations. We suggest that another run of Simulation NS-4 could yield a significantly different average autocorrelation, implying that it is necessary to average more than 32 trials for this simulation to converge to a particular set of values. The curve to which it converges would then likely follow




Figure 4.9: *Simulations 1, NS-1, NS-2 and NS-3 — Average Autocorrelation of 32 Trials*
 $H = 0.75$, $M = \{32, 8, 4, 2\}$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$

Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
1	0.4162	-0.00194	0.0229	0.00666	0.4142	Yes
2	0.2696	0.00001	0.0282	0.00821	0.2696	Yes
4	0.1867	0.00156	0.0299	0.00870	0.1882	Yes
8	0.1305	0.00217	0.0300	0.00872	0.1327	Yes
16	0.0907	0.00304	0.0305	0.00886	0.0938	Yes
32	0.0614	0.00485	0.0303	0.00882	0.0663	Yes
64	0.0431	0.00373	0.0304	0.00885	0.0469	Yes
128	0.0288	0.00439	0.0290	0.00843	0.0331	Yes
256	0.0194	0.00401	0.0278	0.00809	0.0234	Yes
512	0.0135	0.00308	0.0236	0.00687	0.0166	Yes

Table 4.7: *Simulation NS-1 — Autocorrelation Data of 32 Trials*
 $H = 0.75$, $M = 8$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$

Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
1	0.4093	0.00495	0.0284	0.00825	0.4142	Yes
2	0.2604	0.00928	0.0354	0.01030	0.2696	Yes
4	0.1745	0.01378	0.0386	0.01123	0.1882	No
8	0.1186	0.01408	0.0399	0.01160	0.1327	No
16	0.0784	0.01539	0.0420	0.01221	0.0938	No
32	0.0497	0.01664	0.0431	0.01252	0.0663	No
64	0.0309	0.01594	0.0419	0.01217	0.0469	No
128	0.0203	0.01289	0.0395	0.01149	0.0331	No
256	0.0130	0.01040	0.0381	0.01108	0.0234	Yes
512	0.0087	0.00783	0.0334	0.00972	0.0166	Yes

Table 4.8: *Simulation NS-2 — Autocorrelation Data of 32 Trials*
 $H = 0.75$, $M = 4$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$

Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
1	0.4068	0.00741	0.0238	0.00693	0.4142	No
2	0.2576	0.01201	0.0291	0.00847	0.2696	No
4	0.1740	0.01429	0.0318	0.00924	0.1882	No
8	0.1178	0.01493	0.0328	0.00953	0.1327	No
16	0.0773	0.01648	0.0327	0.00951	0.0938	No
32	0.0478	0.01849	0.0285	0.00828	0.0663	No
64	0.0285	0.01839	0.0247	0.00719	0.0469	No
128	0.0137	0.01941	0.0189	0.00550	0.0331	No
256	0.0039	0.01949	0.0158	0.00460	0.0234	No
512	0.0018	0.01476	0.0080	0.00232	0.0166	No

Table 4.9: *Simulation NS-3 — Autocorrelation Data of 32 Trials*
 $H = 0.75$, $M = 2$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$

Sim	σ_{lim}^2	$\hat{\sigma}_{norm}^2$	$\Delta\sigma^2$	$\hat{\sigma}_{\sigma^2}$	$\pm 90\%$ C.I.	Target in C.I.
NS-1	0.2566	0.2534	0.00318	0.0100	0.00292	No
NS-2	0.2566	0.2513	0.00527	0.0135	0.00394	No
NS-3	0.2566	0.2496	0.00701	0.0101	0.00294	No
NS-4	0.2566	0.2635	-0.00693	0.0692	0.02012	Yes

Table 4.10: *Simulations NS-1, NS-2 and NS-3 — Variance Data of 32 Trials*
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$




Figure 4.10: *Simulation NS-4 — Average Autocorrelation of 32 Trials*
 $H = 0.75$, $M = 1$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$




Figure 4.11: *Simulation NS-4 — Autocorrelations for 8 of 32 Individual Trials*
 $H = 0.75$, $M = 1$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$

Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
1	0.4211	-0.00690	0.0780	0.02268	0.4142	Yes
2	0.2749	-0.00525	0.0984	0.02861	0.2696	Yes
4	0.1935	-0.00523	0.1092	0.03174	0.1882	Yes
8	0.1382	-0.00545	0.1154	0.03357	0.1327	Yes
16	0.0984	-0.00460	0.1187	0.03451	0.0938	Yes
32	0.0708	-0.00448	0.1172	0.03409	0.0663	Yes
64	0.0496	-0.00269	0.1145	0.03329	0.0469	Yes
128	0.0346	-0.00148	0.1110	0.03228	0.0331	Yes
256	0.0259	-0.00246	0.1070	0.03113	0.0234	Yes
512	0.0203	-0.00376	0.0985	0.02864	0.0166	Yes

Table 4.11: *Simulation NS-4 — Autocorrelation Data of 32 Trials*
 $H = 0.75$, $M = 1$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$

the same trend as Simulations NS-2 and NS-3, and fall slightly below the target. The variance of Simulation NS-4 is also slightly larger than the target, though the standard deviation of the 32 trials is again large. This creates a wide confidence interval that does contain the target.

It was determined in Section 4.1 that setting $T = 1024 \times 10^4$ in Simulation ST-1 caused the autocorrelations to deviate further from the target by an order of magnitude when compared with Simulation 1 where $T = 1024 \times 10^5$. We now investigate the effects of varying the number of sources while using a simulation time of $T = 1024 \times 10^4$. Simulations NS-5, NS-6, NS-7 and NS-8 were run with 8, 4, 2, and 1 sources, respectively, and the results more closely match what we expected: as the number of sources increases, the correlations converge to the target (Figure 4.12). In all cases, the autocorrelations become essentially zero after a lag of 32–128. See Appendix A for complete autocorrelation and variance data for these trials.

Simulations NS-9 and NS-10 were run with 64 and 128 sources, respectively. The results show that in order to generate traffic with autocorrelations that are as close to the target as Simulation 1, while simulating less time, up to four times as many sources are necessary (see Figure 4.13). This increases simulation clock-time by a factor of four, and reduces the feasibility of using this technique to physically synthesize realistic traffic with a small network of machines.




Figure 4.12: *Simulations ST-1, NS-5, NS-6, NS-7, and NS-8 — Average Autocorrelation of 32 Trials*
 $H = 0.75$, $M = \{32, 8, 4, 2, 1\}$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^4$




Figure 4.13: *Simulations ST-2, NS-9 and NS-10 — Average Autocorrelation of 32 Trials*
 $H = 0.75$, $M = \{32, 64, 128\}$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^4$

4.3 Length of Aggregation Interval

The smoothness of the autocorrelation plots is certainly a result of the number of data-points used in the autocorrelation calculations. If the number of data-points also affects the degree of self-similarity, then one possible way to retain reasonably accurate over a shorter simulation time would be to reduce the interval width, B . This possibility is investigated, but first we explore the effects of varying the aggregation time on the autocorrelations and variances while keeping the other parameters constant.

Simulations AT-1, AT-2, and AT-3 were run with an aggregation time, B , of 512, 128, and 64, respectively. As evident from Figure 4.14 and Tables 4.12, 4.13, and 4.14, values of $B \geq 128$ produce nearly identical results, but at $B = 64$, the results begin to diverge from the target. Simulations AT-4, AT-5, AT-6, and AT-7, run with $B = 32, 16, 4$, and 1 , respectively, follow this trend of increasing correlations (Figure 4.15).

With a smaller interval width, B , there is a higher probability that ON- and OFF-periods will span multiple intervals, resulting in higher correlations. The above results confirm with this observation, and indicate that $B = 128$ is the lowest tested value that results in autocorrelations that are close to the target.

To investigate the effects of B and T together, we run Simulations AT-8, and AT-9 with $B = 128$, $T = 1024 \times 10^4$ and 1024×10^3 , respectively. These results, in Figures 4.16 and 4.17 and Tables 4.15 and 4.16, show Simulation AT-8 to be as close to the target as Simulation 1, and also as close as Simulation AT-2, despite a reduction of T by a factor of ten. This indicates that the number of data-points, $N = \frac{T}{B}$, rather than the total simulation time, T , has the largest effect on the convergence of the autocorrelations to the target. The results of Simulation AT-9 and Simulation ST-2 are again comparable despite a reduction in T by a factor of ten. Simulation AT-9 also shows that with a similar number of data-points, a smaller B results in a higher correlation, consistent with the results of the previous simulations, but amplified by the lower value of T .




Figure 4.14: *Simulations 1, AT-1, AT-2, and AT-3 — Average Autocorrelation of 32 Trials*
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = \{1024, 512, 128, 64\}$, $T = 1024 \times 10^5$

Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
1	0.4249	-0.010671	0.0367	0.010665	0.4142	No
2	0.2795	-0.009879	0.0460	0.013383	0.2696	Yes
4	0.1975	-0.009252	0.0518	0.015058	0.1882	Yes
8	0.1409	-0.008231	0.0551	0.016014	0.1327	Yes
16	0.1016	-0.007859	0.0577	0.016777	0.0938	Yes
32	0.0752	-0.008901	0.0595	0.017291	0.0663	Yes
64	0.0554	-0.008555	0.0602	0.017513	0.0469	Yes
128	0.0414	-0.008300	0.0608	0.017689	0.0331	Yes
256	0.0308	-0.007375	0.0616	0.017918	0.0234	Yes
512	0.0234	-0.006876	0.0617	0.017928	0.0166	Yes

Table 4.12: *Simulation AT-1 — Autocorrelation Data of 32 Trials*
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 512$, $T = 1024 \times 10^5$

Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
1	0.4251	-0.010896	0.0047	0.001362	0.4142	No
2	0.2758	-0.006172	0.0058	0.001694	0.2696	No
4	0.1922	-0.003959	0.0064	0.001873	0.1882	No
8	0.1346	-0.001853	0.0065	0.001898	0.1327	Yes
16	0.0944	-0.000657	0.0068	0.001981	0.0938	Yes
32	0.0660	0.000257	0.0075	0.002184	0.0663	Yes
64	0.0463	0.000538	0.0073	0.002133	0.0469	Yes
128	0.0318	0.001373	0.0070	0.002035	0.0331	Yes
256	0.0221	0.001296	0.0071	0.002057	0.0234	Yes
512	0.0157	0.000878	0.0074	0.002148	0.0166	Yes

Table 4.13: *Simulation AT-2 — Autocorrelation Data of 32 Trials*
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 128$, $T = 1024 \times 10^5$

Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
1	0.4304	-0.016191	0.0034	0.000981	0.4142	No
2	0.2793	-0.009618	0.0041	0.001196	0.2696	No
4	0.1944	-0.006162	0.0047	0.001375	0.1882	No
8	0.1361	-0.003388	0.0048	0.001395	0.1327	No
16	0.0954	-0.001654	0.0050	0.001455	0.0938	No
32	0.0667	-0.000377	0.0051	0.001477	0.0663	No
64	0.0464	0.000513	0.0053	0.001549	0.0469	Yes
128	0.0319	0.001231	0.0050	0.001444	0.0331	Yes
256	0.0225	0.000974	0.0052	0.001507	0.0234	Yes
512	0.0151	0.001521	0.0046	0.001330	0.0166	No

Table 4.14: *Simulation AT-3 — Autocorrelation Data of 32 Trials*
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 64$, $T = 1024 \times 10^5$

Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
1	0.4239	-0.009640	0.0127	0.003706	0.4142	No
2	0.2748	-0.005127	0.0154	0.004482	0.2696	No
4	0.1913	-0.003050	0.0166	0.004820	0.1882	Yes
8	0.1323	0.000364	0.0192	0.005573	0.1327	Yes
16	0.0913	0.002428	0.0201	0.005858	0.0938	Yes
32	0.0637	0.002640	0.0200	0.005828	0.0663	Yes
64	0.0442	0.002644	0.0199	0.005783	0.0469	Yes
128	0.0285	0.004668	0.0194	0.005652	0.0331	Yes
256	0.0175	0.005917	0.0192	0.005573	0.0234	No
512	0.0106	0.005950	0.0175	0.005082	0.0166	No

Table 4.15: *Simulation AT-8 — Autocorrelation Data of 32 Trials*
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 128$, $T = 1024 \times 10^4$




Figure 4.15: Simulations AT-4, AT-5, AT-6, and AT-7 — Average Autocorrelation of 32 Trials
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = \{32, 16, 4, 1\}$, $T = 1024 \times 10^5$




Figure 4.16: Simulations 1 and AT-8 — Average Autocorrelation of 32 Trials
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = \{1024, 128\}$, $T = \{1024 \times 10^5, 1024 \times 10^4\}$




Figure 4.17: *Simulations AT-9 and ST-1 — Average Autocorrelation of 32 Trials*
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = \{128, 1024\}$, $T = \{1024 \times 10^3, 1024 \times 10^4\}$

Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
1	0.4166	-0.002371	0.0277	0.008056	0.4142	Yes
2	0.2651	0.004527	0.0351	0.010198	0.2696	Yes
4	0.1823	0.005992	0.0432	0.012571	0.1882	Yes
8	0.1257	0.006973	0.0406	0.011800	0.1327	Yes
16	0.0819	0.011914	0.0445	0.012934	0.0938	Yes
32	0.0538	-0.053845	0.0449	0.014495	0.0663	No
64	0.0396	-0.039623	0.0343	0.011055	0.0469	No
128	0.0233	-0.023292	0.0298	0.009607	0.0331	No
256	0.0113	-0.011321	0.0313	0.010107	0.0234	No
512	0.0075	-0.007501	0.0270	0.008715	0.0166	Yes

Table 4.16: *Simulation AT-9 — Autocorrelation Data of 32 Trials*
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 128$, $T = 1024 \times 10^3$

4.4 Hurst Parameter

A Hurst parameter of $H = 0.75$ is a good initial simulation choice for two reasons. First, it is approximately equal to the average value of H found by Klivansky et al. [1994] to represent traces of all traffic at NSFNET switches. Second, 0.75 is exactly in the middle of the range $[0.5, 1.0]$, values of H which give rise to long-range dependence in network traffic. Some other mechanisms for approximating fGn, such as *Random Midpoint Displacement* tend to produce higher than target correlations for $H < 0.75$ and lower than target correlations for $H > 0.75$ [Popescu 1999], suggesting that we might see the best results at $H = 0.75$ as well.

We now run simulations with different values of H , roughly corresponding to values of H determined from various network traces. No major studies found H to be significantly higher than 0.9, so we choose that as our upper bound. We use $H = 0.8$ and $H = 0.85$ as intermediate values corresponding to values from traffic examined by Leland et al. [1994]. For values of $H < 0.75$, we choose $H = 0.7$, $H = 0.65$ and $H = 0.60$. Though no major studies examined traffic traces where H was found to be less than 0.69, we choose the lower values for symmetry.

We first notice from Simulation H-1, where $H = 0.80$, that the results are at least as close to the target as Simulation 1 (Figure 4.18 and Table 4.17). However, for Simulation H-2, where $H = 0.85$, the results begin to diverge from the target (Figure 4.19 and Table 4.18). This is even more evident in Simulation H-3, where $H = 0.90$ (Figure 4.20 and Table 4.19). The autocorrelation of Simulation H-3 is less correlated than the target by a significant amount.

A similar phenomenon is observed for values of $H < 0.75$. Simulation H-4, with $H = 0.70$, is close to the target (Figure 4.20 and Table 4.19), Simulation H-5 with $H = 0.65$ and Simulation H-6 with $H = 0.60$ diverge from the target, but in this case are more strongly correlated.

Autocorrelations of simulations with $H > 0.75$ diverge from the target with larger values of the lag, and simulations with $H < 0.75$ converge to the target with a larger lag. However, for $H < 0.75$ the target values approach 0 quite rapidly (i.e. after a lag > 32 for Simulation H-6), implying




Figure 4.18: *Simulation H-1 — Average Autocorrelation of 32 Trials*
 $H = 0.80, M = 32, B = 1024, \mu_{on} = \mu_{off} = 1, T = 1024 \times 10^5$

Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\% \text{ C.I.}$	Target	Target in C.I.
1	0.5173	-0.00162	0.0312	0.00908	0.5157	Yes
2	0.3694	-0.00110	0.0407	0.01183	0.3683	Yes
4	0.2772	-0.00069	0.0467	0.01359	0.2765	Yes
8	0.2095	-0.00046	0.0516	0.01500	0.2091	Yes
16	0.1586	-0.00020	0.0547	0.01592	0.1584	Yes
32	0.1208	-0.00080	0.0574	0.01670	0.1200	Yes
64	0.0914	-0.00045	0.0598	0.01739	0.0909	Yes
128	0.0686	0.00030	0.0595	0.01730	0.0689	Yes
256	0.0504	0.00186	0.0615	0.01789	0.0522	Yes
512	0.0381	0.00149	0.0609	0.01771	0.0396	Yes

Table 4.17: *Simulation H-1 — Autocorrelation Data of 32 Trials*
 $H = 0.80, M = 32, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^5$




Figure 4.19: *Simulation H-2 — Average Autocorrelation of 32 Trials*
 $H = 0.85, M = 32, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^5$

Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\% \text{ C.I.}$	Target	Target in C.I.
1	0.6163	0.008160	0.0180	0.005236	0.6245	No
2	0.4758	0.011710	0.0246	0.007152	0.4875	No
4	0.3795	0.013887	0.0289	0.008396	0.3934	No
8	0.3023	0.016765	0.0318	0.009237	0.3190	No
16	0.2404	0.018666	0.0343	0.009960	0.2590	No
32	0.1895	0.020862	0.0363	0.010567	0.2104	No
64	0.1501	0.020732	0.0381	0.011071	0.1709	No
128	0.1176	0.021166	0.0393	0.011424	0.1388	No
256	0.0912	0.021538	0.0428	0.012441	0.1127	No
512	0.0723	0.019218	0.0417	0.012123	0.0916	No

Table 4.18: *Simulation H-2 — Autocorrelation Data of 32 Trials*
 $H = 0.85, M = 32, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^5$




Figure 4.20: *Simulation H-3 — Average Autocorrelation of 32 Trials*
 $H = 0.90$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$

Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
1	0.6019	0.139225	0.2636	0.076662	0.7411	No
2	0.4980	0.132103	0.2189	0.063667	0.6301	No
4	0.4206	0.125757	0.1860	0.054097	0.5464	No
8	0.3536	0.121523	0.1580	0.045953	0.4752	No
16	0.2947	0.118826	0.1339	0.038929	0.4136	No
32	0.2451	0.114902	0.1142	0.033207	0.3600	No
64	0.1989	0.114456	0.0966	0.028077	0.3134	No
128	0.1602	0.112670	0.0830	0.024130	0.2728	No
256	0.1246	0.112900	0.0741	0.021547	0.2375	No
512	0.0962	0.110604	0.0646	0.018777	0.2068	No

Table 4.19: *Simulation H-3 — Autocorrelation Data of 32 Trials*
 $H = 0.90$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$




Figure 4.21: *Simulation H-4 — Average Autocorrelation of 32 Trials*
 $H = 0.70$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$

Lag	$\hat{\mu}_r(k)$	$\Delta\mu$	$\hat{\sigma}_r(k)$	$\pm 90\%$ C.I.	Target	Target in C.I.
1	0.3278	-0.008337	0.0120	0.003493	0.3195	No
2	0.1934	-0.004688	0.0145	0.004220	0.1888	No
4	0.1249	-0.002426	0.0148	0.004310	0.1225	Yes
8	0.0808	-0.000268	0.0146	0.004249	0.0805	Yes
16	0.0536	-0.000569	0.0139	0.004045	0.0531	Yes
32	0.0335	0.001504	0.0127	0.003697	0.0350	Yes
64	0.0202	0.002897	0.0123	0.003568	0.0231	Yes
128	0.0130	0.002185	0.0123	0.003575	0.0152	Yes
256	0.0081	0.001980	0.0109	0.003176	0.0101	Yes
512	0.0035	0.003130	0.0070	0.002024	0.0066	No

Table 4.20: *Simulation H-4 — Autocorrelation Data of 32 Trials*
 $H = 0.70$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$




Figure 4.22: *Simulation H-5 — Average Autocorrelation of 32 Trials*
 $H = 0.65$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$

Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
1	0.2486	-0.017461	0.0116	0.003370	0.2311	No
2	0.1314	-0.008086	0.0133	0.003878	0.1233	No
4	0.0785	-0.004149	0.0139	0.004048	0.0744	No
8	0.0464	-0.000856	0.0139	0.004047	0.0456	Yes
16	0.0276	0.000363	0.0125	0.003636	0.0280	Yes
32	0.0161	0.001143	0.0119	0.003449	0.0172	Yes
64	0.0090	0.001659	0.0115	0.003350	0.0106	Yes
128	0.0059	0.000647	0.0099	0.002891	0.0065	Yes
256	0.0034	0.000650	0.0081	0.002352	0.0040	Yes
512	0.0022	0.000302	0.0052	0.001519	0.0025	Yes

Table 4.21: *Simulation H-5 — Autocorrelation Data of 32 Trials*
 $H = 0.65$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$




Figure 4.23: *Simulation H-6 — Average Autocorrelation of 32 Trials*
 $H = 0.60$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$

Lag	$\hat{\mu}_r(k)$	$\Delta\mu$	$\hat{\sigma}_r(k)$	$\pm 90\%$ C.I.	Target	Target in C.I.
1	0.1581	-0.009439	0.0612	0.017800	0.1487	Yes
2	0.0756	-0.004380	0.0304	0.008827	0.0712	Yes
4	0.0416	-0.001682	0.0181	0.005272	0.0399	Yes
8	0.0245	-0.001725	0.0124	0.003610	0.0228	Yes
16	0.0143	-0.001217	0.0085	0.002483	0.0131	Yes
32	0.0084	-0.000938	0.0068	0.001988	0.0075	Yes
64	0.0038	0.000500	0.0051	0.001496	0.0043	Yes
128	0.0022	0.000262	0.0048	0.001406	0.0025	Yes
256	0.0011	0.000290	0.0030	0.000868	0.0014	Yes
512	-0.0004	0.001220	0.0031	0.000910	0.0008	No

Table 4.22: *Simulation H-6 — Autocorrelation Data of 32 Trials*
 $H = 0.60$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$

virtually no correlation between elements a distance of 32 apart.

Geist and Westall [2000] provide a statistical analysis of their own procedure for producing samples that are approximately fGn and that of Paxson [1995]. It is interesting to note that the results of those simulations, which target a Hurst parameter of 0.845, roughly match ours for Simulation H-2 in terms of both the distance of the autocorrelations from the target, and the width of the 90% confidence interval.

4.5 Mean ON \neq Mean OFF

The results thus far, where $\mu_{on} = \mu_{off}$, are interesting in a theoretical sense. However, in order to synthesize traffic that is self-similar in nature on a real network, the average link utilization must be kept lower than 100%. There are two ways to achieve this. The first is to specify that on ON-period represents a source sending at some constant rate less than the link bandwidth. The mathematical result presented in [Taqqu et al. 1997] holds for any constant rate. Using this model, we would thus expect our current results to agree with those from a physical network simulation, at least where the transmission rates of the ON-sources are sufficiently low to prevent utilization bursts in excess of 100%. It can be somewhat challenging, however, to produce traffic at a precise rate on a real machine due to characteristics of the physical, data-link, and network layers. Another simple way to achieve a desired link utilization, U , is to set μ_{off} is such a way that $U = M \frac{\mu_{on}}{\mu_{on} + \mu_{off}}$. That is, $\mu_{off} = \frac{M \times \mu_{on}}{U} - \mu_{on}$. This is the scenario we investigate in this section. Again, the simulator does not attempt to capture the effects of short-term utilization bursts that exceed 100%.

For Simulations MO-1, MO-2, MO-3, MO-4, MO-5, and MO-6, the average desired link utilization was set to 25%. For each of these simulations, $\mu_{off} = 4M\mu_{on} - \mu_{on}$. As shown in Figure 4.24, no trend can easily be found among these simulations. This is likely due to the fact that μ_{off} is much closer to B — the value of B is only four times μ_{off} for Simulation MO-1.

To keep the ratio of $\frac{B}{\mu_{on} + \mu_{off}}$ similar to that of Simulation 1, we run Simulations MO-7, MO-8,




Figure 4.24: *Simulations MO- $\{1, 2, 3, 4, 5, 6\}$ — Average Autocorrelation of 32 Trials*
 $H = 0.75$, $M = \{32, 16, 8, 4, 2, 1\}$, $\mu_{on} = 1$, $\mu_{off} = (4M\mu_{on} - \mu_{on})$, $B = 1024$, $T = 1024 \times 10^5$




Figure 4.25: *Simulations MO- $\{7, 8, 9, 10, 11, 12\}$ — Average Autocorrelation of 32 Trials*
 $H = 0.75$, $M = \{32, 16, 8, 4, 2, 1\}$, $\mu_{on} = 1$, $\mu_{off} = (4M\mu_{on} - \mu_{on})$, $B = 16,384$, $T = 16,384 \times 10^5$

MO-9, MO-10, MO-11, and MO-12 with the same values of μ_{on} , μ_{off} , and M as before, but with $B = 16,384$. We increase T to $16,384 \times 10^5$ to maintain 10^5 data-points. The results of these simulations are given in Figure 4.25, and again, no clear trend is observed between the number of sources and the distance of the autocorrelations to the target. If such a trend exists, it will only be determined after a much more thorough examination of the many combinations of simulation parameters.

4.6 Strictly Alternating Sources

Though idealized ON/OFF-sources are considered in [Willinger et al. 1995], later Willinger et al. [1997] and Taqqu et al. [1997] focus on strictly alternating ON/OFF-sources. We now investigate the results of simulations run with strictly alternating sources.

Simulations STR-1, STR-2, STR-3, and STR-4 were run with 32, 8, 4, and 2 sources respectively, and the results are given in Figure 4.26. There is only a slight difference, respectively, between these and Simulations 1, NS-1, NS-2, and NS-3.

In an effort to cause the results to deviate further from the target, Simulations STR-5, STR-6, STR-7, and STR-8 were run with the same respective values of M , but with $T = 1024 \times 10^4$. These values, compared with the corresponding idealized Simulations ST-1, NS-5, NS-6, and NS-6, are closer to the target in every case (except for the tail of Simulation STR-8, which becomes slightly negative). See Appendix A for complete data from these simulations.

Simulations STR-9, STR-10, and STR-11 were run with $H = 0.85, 0.90$, and 0.65 , respectively to compare the results with Simulations H-2, H-3, and H-5. Figures 4.28, 4.29, and 4.30 show that the strictly alternating sources are *not* more strongly correlated and are *not* closer to the target in this case.

Simulation AT-7, with parameters $H = 0.75$, $M = 32$, $B = 1$, $T = 1024 \times 10^5$, and idealized sources, was more strongly correlated than the target by a significant amount. Simulation STR-




Figure 4.26: *Simulations STR-1, STR-2, STR-3, and STR-4 — Average Autocorrelation of 32 Trials Strictly Alternating $H = 0.75$, $M = \{32, 8, 4, 2\}$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$*




Figure 4.27: *Simulations STR-5, STR-6, STR-7, and STR-8 — Average Autocorrelation of 32 Trials Strictly Alternating $H = 0.75$, $M = \{32, 8, 4, 2\}$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^4$*




Figure 4.28: *Simulations STR-9 and H-2 — Average Autocorrelation of 32 Trials* $\{\text{Strictly Alt., Idealized}\}$ $H = 0.85$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$

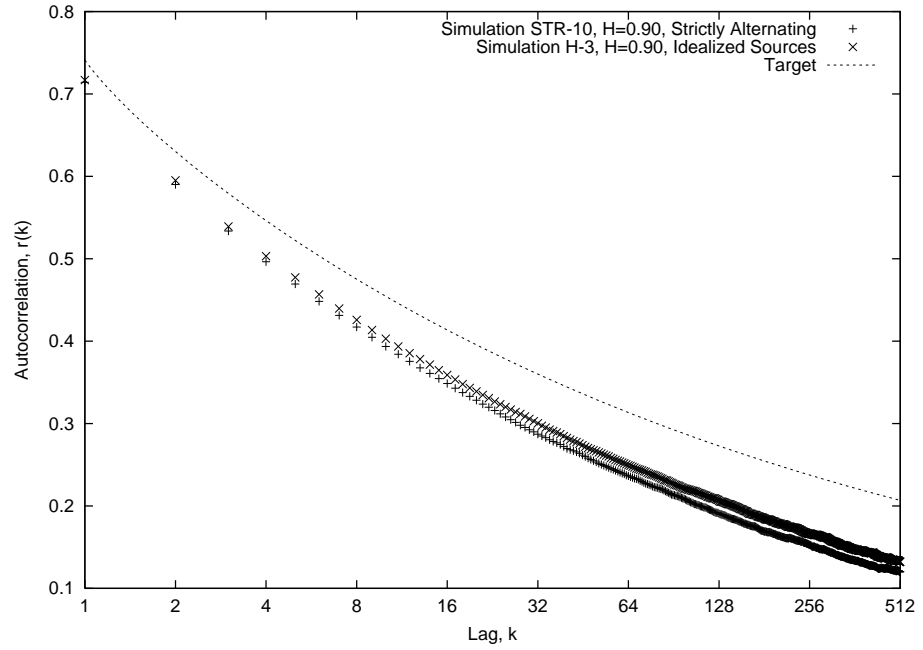


Figure 4.29: *Simulations STR-10 and H-3 — Average Autocorrelation of 32 Trials* $\{\text{Strictly Alt., Idealized}\}$ $H = 0.90$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$

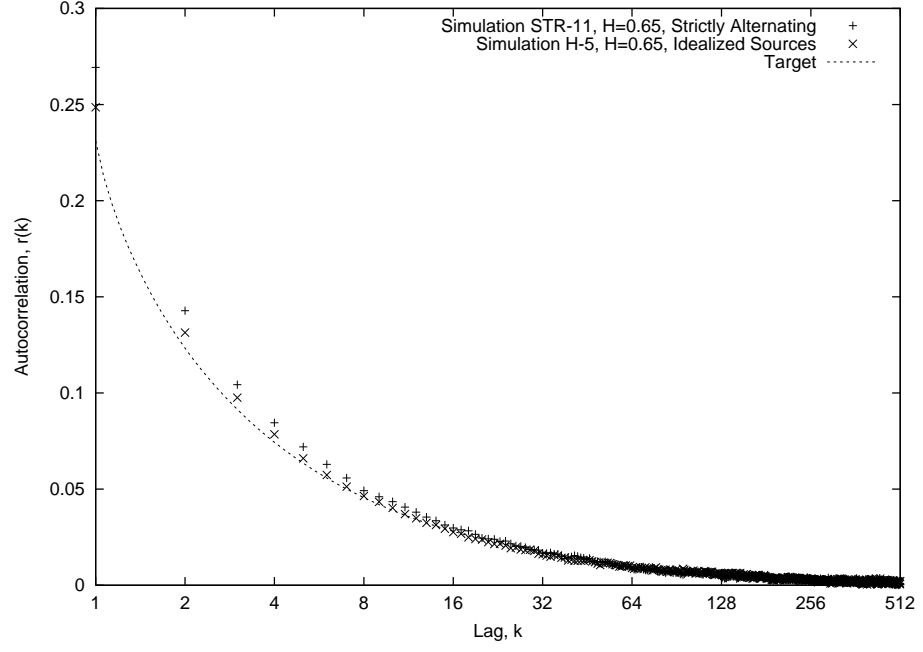


Figure 4.30: *Simulations STR-11 and H-5 — Average Autocorrelation of 32 Trials* $\{\text{Strictly Alt., Idealized}\}$ $H = 0.65$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$




Figure 4.31: *Simulations STR-12 and AT-7 — Average Autocorrelation of 32 Trials* $\{\text{Strictly Alt., Idealized}\}$ $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1$, $T = 1024 \times 10^5$

12 was run with strictly alternating sources, and the results are compared with Simulation AT-7 in Figure 4.31. The results of Simulation STR-12 are again more strongly correlated than the corresponding idealized source simulation. It is interesting that the autocorrelation for a lag of one is the same for both simulations, though the shapes of the curves are different.

It appears that the results of simulations with strictly alternating sources tend to be more strongly correlated than to simulations with idealized sources for an target $H = 0.75$. Strictly alternating simulations with other target values of H are not more strongly correlated, though only a very limited set of simulations were examined.

4.7 Discrete Time

Though a discrete time model is presented in [Willinger et al. 1997] and [Taqqu et al. 1997], the simulations thus far were consistent with a continuous time model, at least to the point allowed by the precision of a 64-bit `double`. To investigate any differences that discretization may cause, the lengths of the ON- and OFF-periods for the simulations in this section have been truncated to a specified discretization value.

The discretization levels for Simulations D-1, D-2, D-3 are set to 0.01, 0.1, and 1.0, respectively. The other parameters are identical to that of Simulation 1, $H = 0.75$, $M = 32$ sources, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$, and idealized ON/OFF-sources. A value of 0.01 is sufficiently smaller than μ_{on} and μ_{off} that we would expect this simulation be be very similar to Simulation 1. Indeed, we see from Figure 4.32 that this is true. It is surprising, however, that a value of 0.1, which is only a tenth of the mean, and a value of 1.0, which is equal to the mean, should also result in nearly identical results. For Simulation D-3, we reason that since every ON or OFF time chosen below the mean is truncated to zero, they have absolutely no effect on the simulation since we are using idealized sources. Thus, the mean of the ON and OFF-periods that are actually used is shifted to a somewhat larger value. This value is apparently not sufficiently closer to B or T to cause a change




Figure 4.32: *Simulations D-1, D-2, and D-3 — Average Autocorrelation of 32 Trials*
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$, Discretized at $\{0.01, 0.1, 1.0\}$

Sim	σ_{lim}^2	$\hat{\sigma}_{norm}^2$	$\Delta\sigma^2$	$\hat{\sigma}_{\sigma^2}$	$\pm 90\%$ C.I.	Target in C.I.
D-1	0.2566	0.2495	0.00707	0.0116	0.00337	No
D-2	0.2566	0.2656	-0.00898	0.0045	0.00131	No
D-3	0.2566	0.4925	-0.23598	0.0096	0.00280	No

Table 4.23: *Simulations D-1, D-2, and D-3 — Average Variance of 32 Trials*
 $H = 0.75$, $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$, Discretized at $\{0.01, 0.1, 1.0\}$

in the results.

Interestingly, the variance from Simulation D-3 is much larger than the target (Table 4.23). Again, we reason that because the smallest ON- and OFF-period for Simulation D-3 is 1.0, the mean ON- and OFF-period of this head-truncated Pareto is likely to be ≈ 3 , since $\frac{1}{3}$ is the ratio of $\frac{\beta}{\mu}$ for a Pareto with $H = 0.75$. The variance appears to agree with this statement, since the expected variance of a simulation with $\mu_{on} = \mu_{off} = 3$ is 0.4444, which is much closer to the average sample variance of 0.4925.




Figure 4.33: *Simulation P-1 — Autocorrelation of 1 Trial*
 $M = 32$, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$, *Poisson ON/OFF-sources*

The similarity of these three simulations compared with the continuous time model used in Simulation 1 might indicate the existence of a continuous time model and proof of self-similar that parallels the discrete model presented in [Taqqu et al. 1997].

4.8 Poisson ON/OFF-Sources

Though not all of the simulations thus far have converged to the target values, the correlations of the simulated packet arrival counts have all been strongly correlated over a wide range of time scales. We now contrast these results with one final simulation. Simulation P-1 was run using ON/OFF-periods with lengths taken from a Poisson process. We set $M = 32$ sources, $\mu_{on} = \mu_{off} = 1$, $B = 1024$, $T = 1024 \times 10^5$, and idealized ON/OFF-sources as in Simulation 1. The target Hurst parameter, H , has no effect on this simulation. As evident from Figure 4.33, Simulation P-1 is completely uncorrelated over an interval of length 1024.

Simulation P-1 provides simulation evidence for that which has already been proven in [Taqqu et al. 1997]: A heavy-tailed distribution for the ON/OFF-periods is the main ingredient necessary for generating strongly correlated network traffic.

4.9 Application of Units

Using the results obtained thus far in Chapter 4, we conclude that a simulation with $H = 0.75$, $M = 8$, $B = 128$, $T = 128 \times 10^4$, $\mu_{on} = \mu_{off} = 1$, and idealized sources would likely provide a good compromise between an accurate approximation of self-similar network traffic and resource requirements. Synthesizing traffic with those parameters on a ATM network, for example, could be done by setting μ_{on} to 100 53-byte ATM cells, which could be sent at a rate of 5 Mbit/sec, keeping the average utilization of a 155 Mbit/sec link at 13%. This causes the mean length of an ON-period (and an OFF-period) to be ≈ 8.1 milliseconds. The traffic would then be aggregated at an interval width of ≈ 1 second, and the simulation would be run for 2.9 hours.

Chapter 5

Conclusions

Though self-similarity or long-range dependence can never be verified for a finite sample, we have shown that the infinite variance, aggregated, ON/OFF-source model presented in [Taqqu et al. 1997] can be used to synthesize network traffic that is statistically consistent with a sample of a stochastic process that is self-similar.

Although an examination of every combination of simulation parameters is impossible, we have shown evidence for the following statements:

- It is possible to produce traffic that is self-similar in nature with many fewer sources than used in previous studies — “good” results were obtained with as few as 8 sources.
- Simulations with strictly alternating sources typically produce results that are more strongly correlated than simulations with idealized ON/OFF-sources, at least for a target H of 0.75.
- The number of data-points, $N = T/B$, is a more important factor than the total simulation time, T , in determining the accuracy of the approximation.
- Reducing the aggregation interval, B , causes the autocorrelation to increase, with the results diverging from the target for values of $B < 128 \times \mu_{on}$. (All simulations that investigated B were run with $\mu_{off} = \mu_{on}$.)
- The correlation of the resulting traffic is exaggerated for values of $H \leq 0.65$ and underestimated for values of $h \geq 0.85$.
- Little difference is observed whether or not artificial discretization is performed.

We have shown that for every combination of H , M , μ_{on} , μ_{off} , B , and T investigated, the correlations of the simulated packet arrival counts are strongly correlated over a wide range of time scales — in stark contrast to a simulation run with a finite variance ON/OFF-model. We have provided insight into the range of values for simulation parameters that produce traffic that is approximately self-similar in nature, and have demonstrated the effects of limited modifications to these parameters. This would allow other researchers to more easily incorporate self-similar traffic models into their network simulations. Finally, we make the claim that the generation of network

traffic that is approximately self-similar or long-range dependent in nature is not only theoretically possible, but feasible with a small number of sources, greatly increasing the ease with which this method can be used to synthesize traffic on a small network of machines.

Chapter 6

Future Work

Future work in this area should include:

- A more thorough examination of simulation results produced with the many possible combinations of parameters
- Synthesis and statistical analysis of traffic generated on a real network using only a few sources
- Determining the effects of self-similar traffic on the behavior of queuing, congestion and other network design models

Appendix A

Complete Simulation Data

Label	Description
Sim	simulation number
H	target Hurst parameter
μ_{on}	mean length of ON-period
μ_{off}	mean length of OFF-period
M	number of sources
B	length of aggregation interval
T	length of total simulation
Strict	(yes) strictly alternating or (no) idealized sources
Discr	discretization value

Table A.1: *Key to Simulation Parameter Labels*

Sim	H	μ_{on}	μ_{off}	M	B	T	Strict	Discr
1	0.75	1	1	32	1024	1024×10^5	No	n/a
ST-1	0.75	1	1	32	1024	1024×10^4	No	n/a
ST-2	0.75	1	1	32	1024	1024×10^3	No	n/a
ST-3	0.75	1	1	32	1024	1024×10^6	No	n/a
NS-1	0.75	1	1	8	1024	1024×10^5	No	n/a
NS-2	0.75	1	1	4	1024	1024×10^5	No	n/a
NS-3	0.75	1	1	2	1024	1024×10^5	No	n/a
NS-4	0.75	1	1	1	1024	1024×10^5	No	n/a
NS-5	0.75	1	1	8	1024	1024×10^4	No	n/a
NS-6	0.75	1	1	4	1024	1024×10^4	No	n/a
NS-7	0.75	1	1	2	1024	1024×10^4	No	n/a
NS-8	0.75	1	1	1	1024	1024×10^4	No	n/a
NS-9	0.75	1	1	64	1024	1024×10^4	No	n/a
NS-10	0.75	1	1	128	1024	1024×10^4	No	n/a
AT-1	0.75	1	1	32	512	1024×10^5	No	n/a
AT-2	0.75	1	1	32	128	1024×10^5	No	n/a
AT-3	0.75	1	1	32	64	1024×10^5	No	n/a
AT-4	0.75	1	1	32	32	1024×10^5	No	n/a
AT-5	0.75	1	1	32	16	1024×10^5	No	n/a

Table A.2: *Simulation Parameters (Continues...)*

Sim	H	μ_{on}	μ_{off}	M	B	T	Strict	Discr
AT-6	0.75	1	1	32	4	1024×10^5	No	n/a
AT-7	0.75	1	1	32	1	1024×10^5	No	n/a
AT-8	0.75	1	1	32	128	1024×10^4	No	n/a
AT-9	0.75	1	1	32	128	1024×10^3	No	n/a
H-1	0.80	1	1	32	1024	1024×10^5	No	n/a
H-2	0.85	1	1	32	1024	1024×10^5	No	n/a
H-3	0.90	1	1	32	1024	1024×10^5	No	n/a
H-4	0.70	1	1	32	1024	1024×10^5	No	n/a
H-5	0.65	1	1	32	1024	1024×10^5	No	n/a
H-6	0.60	1	1	32	1024	1024×10^5	No	n/a
MO-1	0.75	1	127	32	1024	1024×10^5	No	n/a
MO-2	0.75	1	63	16	1024	1024×10^5	No	n/a
MO-3	0.75	1	31	8	1024	1024×10^5	No	n/a
MO-4	0.75	1	15	4	1024	1024×10^5	No	n/a
MO-5	0.75	1	7	2	1024	1024×10^5	No	n/a
MO-6	0.75	1	3	1	1024	1024×10^5	No	n/a
MO-7	0.75	1	127	32	16384	16384×10^5	No	n/a
MO-8	0.75	1	63	16	16384	16384×10^5	No	n/a
MO-9	0.75	1	31	8	16384	16384×10^5	No	n/a
MO-10	0.75	1	15	4	16384	16384×10^5	No	n/a
MO-11	0.75	1	7	2	16384	16384×10^5	No	n/a
STR-1	0.75	1	1	32	1024	1024×10^5	Yes	n/a
STR-2	0.75	1	1	8	1024	1024×10^5	Yes	n/a
STR-3	0.75	1	1	4	1024	1024×10^5	Yes	n/a
STR-4	0.75	1	1	2	1024	1024×10^5	Yes	n/a
STR-5	0.75	1	1	32	1024	1024×10^4	Yes	n/a
STR-6	0.75	1	1	8	1024	1024×10^4	Yes	n/a
STR-7	0.75	1	1	4	1024	1024×10^4	Yes	n/a
STR-8	0.75	1	1	2	1024	1024×10^4	Yes	n/a
STR-9	0.85	1	1	32	1024	1024×10^5	Yes	n/a
STR-10	0.90	1	1	32	1024	1024×10^5	Yes	n/a
STR-11	0.65	1	1	32	1024	1024×10^5	Yes	n/a
STR-12	0.75	1	1	32	1	1024×10^5	Yes	n/a
D-1	0.75	1	1	32	1024	1024×10^5	Yes	0.01
D-2	0.75	1	1	32	1024	1024×10^5	Yes	0.1
D-3	0.75	1	1	32	1024	1024×10^5	Yes	1.0

Table A.2 Continued: *Simulation Parameters*

Label	Description
Sim	simulation number
σ_{lim}^2	expected variance
$\hat{\sigma}_{norm}^2$	mean sample variance
$\Delta\sigma^2$	difference between $\hat{\sigma}_{norm}^2$ and σ_{lim}^2
$\hat{\sigma}_{\sigma^2}$	standard deviation of the sample variances
$\pm 90\%$ C.I.	half-width of the two-sided 90% confidence interval
Target in C.I.	is the target contained within the 90% confidence interval?

Table A.3: *Key to Variance Data Labels*

Sim	σ_{lim}^2	$\hat{\sigma}_{norm}^2$	$\Delta\sigma^2$	$\hat{\sigma}_{\sigma^2}$	$\pm 90\%$ C.I.	Target in C.I.
1	0.2566	0.2527	0.00392	0.0050	0.00147	No
ST-1	0.2566	0.2492	0.00738	0.0104	0.00304	No
ST-2	0.2566	0.2438	0.01282	0.0309	0.00898	No
ST-3	0.2566	0.2538	0.00279	0.0025	0.00169	No
NS-1	0.2566	0.2534	0.00318	0.0100	0.00292	No
NS-2	0.2566	0.2513	0.00527	0.0135	0.00394	No
NS-3	0.2566	0.2496	0.00701	0.0101	0.00294	No
NS-4	0.2566	0.2635	-0.00693	0.0692	0.02012	Yes
NS-5	0.2566	0.2434	0.01322	0.0136	0.00394	No
NS-6	0.2566	0.2459	0.01072	0.0216	0.00628	No
NS-7	0.2566	0.2431	0.01350	0.0295	0.00857	No
NS-8	0.2566	0.2377	0.01891	0.0308	0.00895	No
NS-9	0.2566	0.2570	-0.00041	0.0193	0.00563	Yes
NS-10	0.2566	0.2515	0.00512	0.0095	0.00277	No
AT-1	0.2566	0.2561	0.00046	0.0225	0.00654	Yes
AT-2	0.2566	0.2488	0.00782	0.0018	0.00051	No
AT-3	0.2566	0.2456	0.01096	0.0013	0.00037	No
AT-4	0.2566	0.2418	0.01479	0.0015	0.00044	No
AT-5	0.2566	0.2357	0.02094	0.0004	0.00012	No
AT-6	0.2566	0.2152	0.04138	0.0002	0.00007	No
AT-7	0.2566	0.1764	0.08024	0.0001	0.00002	No
AT-8	0.2566	0.2484	0.00819	0.0055	0.00161	No
AT-9	0.2566	0.2454	0.01121	0.0116	0.00339	No
H-1	0.2254	0.2255	-0.00006	0.0172	0.00501	Yes
H-2	0.2082	0.2041	0.00412	0.0088	0.00257	No
H-3	0.2022	0.1887	0.01353	0.0127	0.00370	No
H-4	0.3098	0.3001	0.00966	0.0052	0.00150	No
H-5	0.4052	0.3753	0.02989	0.0058	0.00168	No
H-6	0.6050	0.4941	0.11083	0.0061	0.00177	No

Table A.4: *Complete Variance Data (Continues...)*

Sim	σ_{lim}^2	$\hat{\sigma}_{norm}^2$	$\Delta\sigma^2$	$\hat{\sigma}_{\sigma^2}$	$\pm 90\%$ C.I.	Target in C.I.
MO-1	0.0086	0.0087	-0.00007	0.0014	0.00039	Yes
MO-2	0.0175	0.0199	-0.00238	0.0072	0.00209	No
MO-3	0.0355	0.0343	0.00123	0.0027	0.00080	No
MO-4	0.0709	0.0689	0.00207	0.0090	0.00262	Yes
MO-5	0.1354	0.1393	-0.00399	0.0446	0.01296	Yes
MO-6	0.2277	0.3009	-0.07325	0.3886	0.11300	Yes
MO-7	0.0086	0.0093	-0.00069	0.0028	0.00080	Yes
MO-8	0.0175	0.0159	0.00155	0.0022	0.00064	No
MO-9	0.0355	0.0354	0.00008	0.0054	0.00156	Yes
MO-10	0.0709	0.0637	0.00723	0.0057	0.00166	No
MO-11	0.1354	0.1366	-0.00128	0.0312	0.00908	Yes
STR-1	0.2566	0.2563	0.00031	0.0456	0.01326	Yes
STR-2	0.2566	0.2445	0.01209	0.0095	0.00275	No
STR-3	0.2566	0.2440	0.01260	0.0186	0.00542	No
STR-4	0.2566	0.2420	0.01457	0.0123	0.00358	No
STR-5	0.2566	0.2425	0.01410	0.0120	0.00349	No
STR-6	0.2566	0.2488	0.00777	0.0299	0.00870	Yes
STR-7	0.2566	0.2467	0.00995	0.0776	0.02256	Yes
STR-8	0.2566	0.2392	0.01738	0.0277	0.00805	No
STR-9	0.2082	0.2017	0.00645	0.0098	0.00284	No
STR-10	0.2022	0.1855	0.01669	0.0136	0.00396	No
STR-11	0.4052	0.3442	0.06099	0.0061	0.00179	No
STR-12	0.2566	0.1214	0.13516	0.0002	0.00007	No
D-1	0.2566	0.2495	0.00707	0.0116	0.00337	No
D-2	0.2566	0.2656	-0.00898	0.0045	0.00131	No
D-3	0.2566	0.4925	-0.23589	0.0096	0.00280	No

Table A.4 Continued: *Complete Variance Data*

Label	Description
Sim	simulation number
Lag	distance between elements for which the correlation is determined
$\hat{\mu}_{r(k)}$	mean sample autocorrelation
$\Delta\mu$	difference between $\hat{\mu}_{r(k)}$ and the expected value
$\hat{\sigma}_{r(k)}$	standard deviation of the sample autocorrelation values for a given lag
$\pm 90\%$ C.I.	half-width of the two-sided 90% confidence interval
Target	expected value
Target in C.I.	is the target contained within the 90% confidence interval?

Table A.5: *Key to Autocorrelation Data Labels*

Sim	Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
1	$H = 0.75, M = 32, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^5$						
	1	0.4149	-0.00069	0.0125	0.00364	0.4142	Yes
	2	0.2668	0.00288	0.0156	0.00455	0.2696	Yes
	4	0.1841	0.00412	0.0171	0.00499	0.1882	Yes
	8	0.1284	0.00427	0.0176	0.00511	0.1327	Yes
	16	0.0890	0.00480	0.0186	0.00542	0.0938	Yes
	32	0.0614	0.00489	0.0179	0.00520	0.0663	Yes
	64	0.0431	0.00378	0.0188	0.00548	0.0469	Yes
	128	0.0304	0.00277	0.0195	0.00568	0.0331	Yes
	256	0.0195	0.00396	0.0176	0.00513	0.0234	Yes
	512	0.0111	0.00546	0.0175	0.00510	0.0166	No
ST-1	$H = 0.75, M = 32, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^4$						
	1	0.4060	0.008173	0.0248	0.007204	0.4142	No
	2	0.2553	0.014318	0.0269	0.007828	0.2696	No
	4	0.1712	0.017025	0.0351	0.010200	0.1882	No
	8	0.1060	0.026719	0.0355	0.010319	0.1327	No
	16	0.0725	0.021254	0.0355	0.010320	0.0938	No
	32	0.0424	0.023892	0.0302	0.008772	0.0663	No
	64	0.0210	0.025835	0.0292	0.008483	0.0469	No
	128	0.0129	0.020291	0.0244	0.007098	0.0331	No
	256	0.0024	0.020989	0.0211	0.006144	0.0234	No
	512	0.0028	0.013723	0.0159	0.004621	0.0166	No
ST-2	$H = 0.75, M = 32, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^3$						
	1	0.3826	0.031644	0.0668	0.019439	0.4142	No
	2	0.2186	0.051057	0.0778	0.022615	0.2696	No
	4	0.1405	0.047796	0.0832	0.024208	0.1882	No
	8	0.0851	0.047612	0.0839	0.024392	0.1327	No
	16	0.0537	0.040033	0.0933	0.027126	0.0938	No
	32	0.0331	0.033182	0.0874	0.025413	0.0663	No
	64	0.0095	0.037400	0.0810	0.023558	0.0469	No
	128	0.0147	0.018430	0.0819	0.023807	0.0331	Yes
	256	0.0148	0.008640	0.0744	0.021635	0.0234	Yes

Table A.6: *Complete Autocorrelation Data (Continues...)*

Sim	Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
ST-3	$H = 0.75, M = 32, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^6$						
	1	0.4180	-0.00375	0.0074	0.00548	0.4142	Yes
	2	0.2717	-0.00202	0.0086	0.00635	0.2696	Yes
	4	0.1893	-0.00109	0.0092	0.00677	0.1882	Yes
	8	0.1343	-0.00161	0.0096	0.00703	0.1327	Yes
	16	0.0951	-0.00135	0.0096	0.00710	0.0938	Yes
	32	0.0675	-0.00124	0.0107	0.00785	0.0663	Yes
	64	0.0477	-0.00085	0.0099	0.00726	0.0469	Yes
	128	0.0335	-0.00031	0.0105	0.00769	0.0331	Yes
	256	0.0228	0.00065	0.0093	0.00688	0.0234	Yes
	512	0.0149	0.00166	0.0088	0.00650	0.0166	Yes
NS-1	$H = 0.75, M = 8, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^5$						
	1	0.4162	-0.00194	0.0229	0.00666	0.4142	Yes
	2	0.2696	0.00001	0.0282	0.00821	0.2696	Yes
	4	0.1867	0.00156	0.0299	0.00870	0.1882	Yes
	8	0.1305	0.00217	0.0300	0.00872	0.1327	Yes
	16	0.0907	0.00304	0.0305	0.00886	0.0938	Yes
	32	0.0614	0.00485	0.0303	0.00882	0.0663	Yes
	64	0.0431	0.00373	0.0304	0.00885	0.0469	Yes
	128	0.0288	0.00439	0.0290	0.00843	0.0331	Yes
	256	0.0194	0.00401	0.0278	0.00809	0.0234	Yes
	512	0.0135	0.00308	0.0236	0.00687	0.0166	Yes
NS-2	$H = 0.75, M = 4, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^5$						
	1	0.4093	0.00495	0.0284	0.00825	0.4142	Yes
	2	0.2604	0.00928	0.0354	0.01030	0.2696	Yes
	4	0.1745	0.01378	0.0386	0.01123	0.1882	No
	8	0.1186	0.01408	0.0399	0.01160	0.1327	No
	16	0.0784	0.01539	0.0420	0.01221	0.0938	No
	32	0.0497	0.01664	0.0431	0.01252	0.0663	No
	64	0.0309	0.01594	0.0419	0.01217	0.0469	No
	128	0.0203	0.01289	0.0395	0.01149	0.0331	No
	256	0.0130	0.01040	0.0381	0.01108	0.0234	Yes
	512	0.0087	0.00783	0.0334	0.00972	0.0166	Yes
NS-3	$H = 0.75, M = 2, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^5$						
	1	0.4068	0.00741	0.0238	0.00693	0.4142	No
	2	0.2576	0.01201	0.0291	0.00847	0.2696	No
	4	0.1740	0.01429	0.0318	0.00924	0.1882	No
	8	0.1178	0.01493	0.0328	0.00953	0.1327	No
	16	0.0773	0.01648	0.0327	0.00951	0.0938	No
	32	0.0478	0.01849	0.0285	0.00828	0.0663	No
	64	0.0285	0.01839	0.0247	0.00719	0.0469	No
	128	0.0137	0.01941	0.0189	0.00550	0.0331	No
	256	0.0039	0.01949	0.0158	0.00460	0.0234	No
	512	0.0018	0.01476	0.0080	0.00232	0.0166	No

Table A.6 Continued: *Complete Autocorrelation Data*

Sim	Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
NS-4	$H = 0.75, M = 1, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^5$						
	1	0.4211	-0.00690	0.0780	0.02268	0.4142	Yes
	2	0.2749	-0.00525	0.0984	0.02861	0.2696	Yes
	4	0.1935	-0.00523	0.1092	0.03174	0.1882	Yes
	8	0.1382	-0.00545	0.1154	0.03357	0.1327	Yes
	16	0.0984	-0.00460	0.1187	0.03451	0.0938	Yes
	32	0.0708	-0.00448	0.1172	0.03409	0.0663	Yes
	64	0.0496	-0.00269	0.1145	0.03329	0.0469	Yes
	128	0.0346	-0.00148	0.1110	0.03228	0.0331	Yes
	256	0.0259	-0.00246	0.1070	0.03113	0.0234	Yes
	512	0.0203	-0.00376	0.0985	0.02864	0.0166	Yes
NS-5	$H = 0.75, M = 8, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^4$						
	1	0.3898	0.024450	0.0333	0.009697	0.4142	No
	2	0.2365	0.033113	0.0395	0.011485	0.2696	No
	4	0.1519	0.036315	0.0417	0.012125	0.1882	No
	8	0.0918	0.040898	0.0411	0.011953	0.1327	No
	16	0.0502	0.043580	0.0426	0.012382	0.0938	No
	32	0.0305	0.035782	0.0359	0.010429	0.0663	No
	64	0.0175	0.029336	0.0307	0.008936	0.0469	No
	128	0.0055	0.027639	0.0243	0.007071	0.0331	No
	256	0.0045	0.018978	0.0137	0.003975	0.0234	No
	512	-0.0027	0.019267	0.0136	0.003944	0.0166	No
NS-6	$H = 0.75, M = 4, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^4$						
	1	0.3931	0.021141	0.0484	0.014084	0.4142	No
	2	0.2443	0.025344	0.0568	0.016527	0.2696	No
	4	0.1602	0.028015	0.0645	0.018747	0.1882	No
	8	0.1055	0.027187	0.0645	0.018768	0.1327	No
	16	0.0676	0.026222	0.0598	0.017391	0.0938	No
	32	0.0426	0.023677	0.0517	0.015031	0.0663	No
	64	0.0216	0.025273	0.0398	0.011562	0.0469	No
	128	0.0070	0.026164	0.0158	0.004594	0.0331	No
	256	-0.0032	0.026600	0.0116	0.003367	0.0234	No
	512	0.0015	0.015118	0.0171	0.004977	0.0166	No
NS-7	$H = 0.75, M = 2, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^4$						
	1	0.3751	0.039121	0.0563	0.016359	0.4142	No
	2	0.2205	0.049115	0.0683	0.019848	0.2696	No
	4	0.1285	0.059771	0.0695	0.020204	0.1882	No
	8	0.0726	0.060146	0.0685	0.019909	0.1327	No
	16	0.0374	0.056419	0.0618	0.017962	0.0938	No
	32	0.0130	0.053307	0.0571	0.016598	0.0663	No
	64	0.0083	0.038580	0.0483	0.014035	0.0469	No
	128	0.0043	0.028866	0.0347	0.010080	0.0331	No
	256	0.0019	0.021562	0.0139	0.004055	0.0234	No
	512	0.0024	0.014140	0.0134	0.003896	0.0166	No

Table A.6 Continued: *Complete Autocorrelation Data*

Sim	Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
NS-8	$H = 0.75, M = 1, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^4$						
	1	0.3653	0.048876	0.0689	0.020029	0.4142	No
	2	0.2052	0.064457	0.0820	0.023842	0.2696	No
	4	0.1235	0.064789	0.0800	0.023265	0.1882	No
	8	0.0728	0.059902	0.0750	0.021800	0.1327	No
	16	0.0357	0.058052	0.0604	0.017573	0.0938	No
	32	0.0155	0.050838	0.0412	0.011977	0.0663	No
	64	0.0075	0.039368	0.0237	0.006892	0.0469	No
	128	0.0019	0.031250	0.0129	0.003751	0.0331	No
	256	0.0026	0.020877	0.0115	0.003346	0.0234	No
	512	0.0022	0.014354	0.0122	0.003556	0.0166	No
NS-9	$H = 0.75, M = 64, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^4$						
	1	0.4225	-0.008284	0.0396	0.011510	0.4142	Yes
	2	0.2790	-0.009342	0.0489	0.014222	0.2696	Yes
	4	0.1942	-0.005947	0.0557	0.016199	0.1882	Yes
	8	0.1396	-0.006897	0.0601	0.017479	0.1327	Yes
	16	0.1018	-0.007989	0.0645	0.018763	0.0938	Yes
	32	0.0734	-0.007099	0.0622	0.018092	0.0663	Yes
	64	0.0526	-0.005709	0.0631	0.018355	0.0469	Yes
	128	0.0367	-0.003586	0.0601	0.017471	0.0331	Yes
	256	0.0243	-0.000847	0.0561	0.016308	0.0234	Yes
	512	0.0235	-0.006918	0.0550	0.016006	0.0166	Yes
NS-10	$H = 0.75, M = 128, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^4$						
	1	0.4105	0.003685	0.0214	0.006232	0.4142	Yes
	2	0.2643	0.005355	0.0240	0.006976	0.2696	Yes
	4	0.1807	0.007589	0.0216	0.006283	0.1882	No
	8	0.1275	0.005217	0.0252	0.007315	0.1327	Yes
	16	0.0909	0.002832	0.0283	0.008238	0.0938	Yes
	32	0.0605	0.005772	0.0297	0.008651	0.0663	Yes
	64	0.0423	0.004529	0.0277	0.008068	0.0469	Yes
	128	0.0291	0.004011	0.0256	0.007457	0.0331	Yes
	256	0.0178	0.005603	0.0246	0.007162	0.0234	Yes
	512	0.0127	0.003889	0.0230	0.006684	0.0166	Yes
AT-1	$H = 0.75, M = 32, \mu_{on} = \mu_{off} = 1, B = 512, T = 1024 \times 10^5$						
	1	0.4249	-0.010671	0.0367	0.010665	0.4142	No
	2	0.2795	-0.009879	0.0460	0.013383	0.2696	Yes
	4	0.1975	-0.009252	0.0518	0.015058	0.1882	Yes
	8	0.1409	-0.008231	0.0551	0.016014	0.1327	Yes
	16	0.1016	-0.007859	0.0577	0.016777	0.0938	Yes
	32	0.0752	-0.008901	0.0595	0.017291	0.0663	Yes
	64	0.0554	-0.008555	0.0602	0.017513	0.0469	Yes
	128	0.0414	-0.008300	0.0608	0.017689	0.0331	Yes
	256	0.0308	-0.007375	0.0616	0.017918	0.0234	Yes
	512	0.0234	-0.006876	0.0617	0.017928	0.0166	Yes

Table A.6 Continued: *Complete Autocorrelation Data*

Sim	Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
AT-2	$H = 0.75, M = 32, \mu_{on} = \mu_{off} = 1, B = 128, T = 1024 \times 10^5$						
	1	0.4251	-0.010896	0.0047	0.001362	0.4142	No
	2	0.2758	-0.006172	0.0058	0.001694	0.2696	No
	4	0.1922	-0.003959	0.0064	0.001873	0.1882	No
	8	0.1346	-0.001853	0.0065	0.001898	0.1327	Yes
	16	0.0944	-0.000657	0.0068	0.001981	0.0938	Yes
	32	0.0660	0.000257	0.0075	0.002184	0.0663	Yes
	64	0.0463	0.000538	0.0073	0.002133	0.0469	Yes
	128	0.0318	0.001373	0.0070	0.002035	0.0331	Yes
	256	0.0221	0.001296	0.0071	0.002057	0.0234	Yes
	512	0.0157	0.000878	0.0074	0.002148	0.0166	Yes
AT-3	$H = 0.75, M = 32, \mu_{on} = \mu_{off} = 1, B = 64, T = 1024 \times 10^5$						
	1	0.4304	-0.016191	0.0034	0.000981	0.4142	No
	2	0.2793	-0.009618	0.0041	0.001196	0.2696	No
	4	0.1944	-0.006162	0.0047	0.001375	0.1882	No
	8	0.1361	-0.003388	0.0048	0.001395	0.1327	No
	16	0.0954	-0.001654	0.0050	0.001455	0.0938	No
	32	0.0667	-0.000377	0.0051	0.001477	0.0663	No
	64	0.0464	0.000513	0.0053	0.001549	0.0469	Yes
	128	0.0319	0.001231	0.0050	0.001444	0.0331	Yes
	256	0.0225	0.000974	0.0052	0.001507	0.0234	Yes
	512	0.0151	0.001521	0.0046	0.001330	0.0166	No
AT-4	$H = 0.75, M = 32, \mu_{on} = \mu_{off} = 1, B = 32, T = 1024 \times 10^5$						
	1	0.4394	-0.025214	0.0043	0.001237	0.4142	No
	2	0.2862	-0.016534	0.0054	0.001582	0.2696	No
	4	0.1997	-0.011417	0.0060	0.001753	0.1882	No
	8	0.1409	-0.008148	0.0063	0.001833	0.1327	No
	16	0.0996	-0.005852	0.0067	0.001954	0.0938	No
	32	0.0707	-0.004359	0.0069	0.001994	0.0663	No
	64	0.0500	-0.003155	0.0068	0.001970	0.0469	No
	128	0.0355	-0.002317	0.0071	0.002078	0.0331	No
	256	0.0248	-0.001353	0.0073	0.002131	0.0234	Yes
	512	0.0180	-0.001401	0.0075	0.002174	0.0166	Yes
AT-5	$H = 0.75, M = 32, \mu_{on} = \mu_{off} = 1, B = 16, T = 1024 \times 10^5$						
	1	0.4502	-0.035972	0.0012	0.000344	0.4142	No
	2	0.2928	-0.023189	0.0015	0.000444	0.2696	No
	4	0.2042	-0.015905	0.0016	0.000468	0.1882	No
	8	0.1437	-0.010946	0.0016	0.000465	0.1327	No
	16	0.1013	-0.007495	0.0018	0.000520	0.0938	No
	32	0.0712	-0.004870	0.0019	0.000548	0.0663	No
	64	0.0499	-0.003065	0.0018	0.000526	0.0469	No
	128	0.0350	-0.001854	0.0017	0.000505	0.0331	No
	256	0.0245	-0.001028	0.0019	0.000550	0.0234	No
	512	0.0169	-0.000322	0.0017	0.000487	0.0166	Yes

Table A.6 Continued: *Complete Autocorrelation Data*

Sim	Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
AT-6	$H = 0.75, M = 32, \mu_{on} = \mu_{off} = 1, B = 4, T = 1024 \times 10^5$						
	1	0.4924	-0.078170	0.0008	0.000227	0.4142	No
	2	0.3209	-0.051238	0.0010	0.000299	0.2696	No
	4	0.2238	-0.035559	0.0012	0.000354	0.1882	No
	8	0.1575	-0.024830	0.0014	0.000395	0.1327	No
	16	0.1112	-0.017383	0.0014	0.000418	0.0938	No
	32	0.0784	-0.012112	0.0014	0.000401	0.0663	No
	64	0.0551	-0.008267	0.0014	0.000396	0.0469	No
	128	0.0388	-0.005659	0.0014	0.000400	0.0331	No
	256	0.0273	-0.003858	0.0015	0.000426	0.0234	No
	512	0.0190	-0.002420	0.0015	0.000432	0.0166	No
AT-7	$H = 0.75, M = 32, \mu_{on} = \mu_{off} = 1, B = 1, T = 1024 \times 10^5$						
	1	0.5940	-0.179768	0.0005	0.000137	0.4142	No
	2	0.3924	-0.122752	0.0007	0.000200	0.2696	No
	4	0.2740	-0.085725	0.0008	0.000239	0.1882	No
	8	0.1932	-0.060471	0.0009	0.000265	0.1327	No
	16	0.1365	-0.042749	0.0010	0.000289	0.0938	No
	32	0.0966	-0.030266	0.0010	0.000304	0.0663	No
	64	0.0682	-0.021374	0.0010	0.000302	0.0469	No
	128	0.0483	-0.015155	0.0011	0.000309	0.0331	No
	256	0.0342	-0.010748	0.0011	0.000330	0.0234	No
	512	0.0032	-0.007585	0.0011	0.000323	0.0166	No
AT-8	$H = 0.75, M = 32, \mu_{on} = \mu_{off} = 1, B = 128, T = 1024 \times 10^4$						
	1	0.4239	-0.009640	0.0127	0.003706	0.4142	No
	2	0.2748	-0.005127	0.0154	0.004482	0.2696	No
	4	0.1913	-0.003050	0.0166	0.004820	0.1882	Yes
	8	0.1323	0.000364	0.0192	0.005573	0.1327	Yes
	16	0.0913	0.002428	0.0201	0.005858	0.0938	Yes
	32	0.0637	0.002640	0.0200	0.005828	0.0663	Yes
	64	0.0442	0.002644	0.0199	0.005783	0.0469	Yes
	128	0.0285	0.004668	0.0194	0.005652	0.0331	Yes
	256	0.0175	0.005917	0.0192	0.005573	0.0234	No
	512	0.0106	0.005950	0.0175	0.005082	0.0166	No
AT-9	$H = 0.75, M = 32, \mu_{on} = \mu_{off} = 1, B = 128, T = 1024 \times 10^3$						
	1	0.4166	-0.002371	0.0277	0.008056	0.4142	Yes
	2	0.2651	0.004527	0.0351	0.010198	0.2696	Yes
	4	0.1823	0.005992	0.0432	0.012571	0.1882	Yes
	8	0.1257	0.006973	0.0406	0.011800	0.1327	Yes
	16	0.0819	0.011914	0.0445	0.012934	0.0938	Yes
	32	0.0538	-0.053845	0.0449	0.014495	0.0663	No
	64	0.0396	-0.039623	0.0343	0.011055	0.0469	No
	128	0.0233	-0.023292	0.0298	0.009607	0.0331	No
	256	0.0113	-0.011321	0.0313	0.010107	0.0234	No
	512	0.0075	-0.007501	0.0270	0.008715	0.0166	Yes

Table A.6 Continued: *Complete Autocorrelation Data*

Sim	Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
H-1	$H = 0.80, M = 32, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^5$						
	1	0.5173	-0.00162	0.0312	0.00908	0.5157	Yes
	2	0.3694	-0.00110	0.0407	0.01183	0.3683	Yes
	4	0.2772	-0.00069	0.0467	0.01359	0.2765	Yes
	8	0.2095	-0.00046	0.0516	0.01500	0.2091	Yes
	16	0.1586	-0.00020	0.0547	0.01592	0.1584	Yes
	32	0.1208	-0.00080	0.0574	0.01670	0.1200	Yes
	64	0.0914	-0.00045	0.0598	0.01739	0.0909	Yes
	128	0.0686	0.00030	0.0595	0.01730	0.0689	Yes
	256	0.0504	0.00186	0.0615	0.01789	0.0522	Yes
	512	0.0381	0.00149	0.0609	0.01771	0.0396	Yes
H-2	$H = 0.85, M = 32, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^5$						
	1	0.6163	0.008160	0.0180	0.005236	0.6245	No
	2	0.4758	0.011710	0.0246	0.007152	0.4875	No
	4	0.3795	0.013887	0.0289	0.008396	0.3934	No
	8	0.3023	0.016765	0.0318	0.009237	0.3190	No
	16	0.2404	0.018666	0.0343	0.009960	0.2590	No
	32	0.1895	0.020862	0.0363	0.010567	0.2104	No
	64	0.1501	0.020732	0.0381	0.011071	0.1709	No
	128	0.1176	0.021166	0.0393	0.011424	0.1388	No
	256	0.0912	0.021538	0.0428	0.012441	0.1127	No
	512	0.0723	0.019218	0.0417	0.012123	0.0916	No
H-3	$H = 0.90, M = 32, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^5$						
	1	0.6019	0.139225	0.2636	0.076662	0.7411	No
	2	0.4980	0.132103	0.2189	0.063667	0.6301	No
	4	0.4206	0.125757	0.1860	0.054097	0.5464	No
	8	0.3536	0.121523	0.1580	0.045953	0.4752	No
	16	0.2947	0.118826	0.1339	0.038929	0.4136	No
	32	0.2451	0.114902	0.1142	0.033207	0.3600	No
	64	0.1989	0.114456	0.0966	0.028077	0.3134	No
	128	0.1602	0.112670	0.0830	0.024130	0.2728	No
	256	0.1246	0.112900	0.0741	0.021547	0.2375	No
	512	0.0962	0.110604	0.0646	0.018777	0.2068	No
H-4	$H = 0.70, M = 32, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^5$						
	1	0.3278	-0.008337	0.0120	0.003493	0.3195	No
	2	0.1934	-0.004688	0.0145	0.004220	0.1888	No
	4	0.1249	-0.002426	0.0148	0.004310	0.1225	Yes
	8	0.0808	-0.000268	0.0146	0.004249	0.0805	Yes
	16	0.0536	-0.000569	0.0139	0.004045	0.0531	Yes
	32	0.0335	0.001504	0.0127	0.003697	0.0350	Yes
	64	0.0202	0.002897	0.0123	0.003568	0.0231	Yes
	128	0.0130	0.002185	0.0123	0.003575	0.0152	Yes
	256	0.0081	0.001980	0.0109	0.003176	0.0101	Yes
	512	0.0035	0.003130	0.0070	0.002024	0.0066	No

Table A.6 Continued: *Complete Autocorrelation Data*

Sim	Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
H-5	$H = 0.65, M = 32, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^5$						
	1	0.2486	-0.017461	0.0116	0.003370	0.2311	No
	2	0.1314	-0.008086	0.0133	0.003878	0.1233	No
	4	0.0785	-0.004149	0.0139	0.004048	0.0744	No
	8	0.0464	-0.000856	0.0139	0.004047	0.0456	Yes
	16	0.0276	0.000363	0.0125	0.003636	0.0280	Yes
	32	0.0161	0.001143	0.0119	0.003449	0.0172	Yes
	64	0.0090	0.001659	0.0115	0.003350	0.0106	Yes
	128	0.0059	0.000647	0.0099	0.002891	0.0065	Yes
	256	0.0034	0.000650	0.0081	0.002352	0.0040	Yes
	512	0.0022	0.000302	0.0052	0.001519	0.0025	Yes
H-6	$H = 0.60, M = 32, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^5$						
	1	0.1581	-0.009439	0.0612	0.017800	0.1487	Yes
	2	0.0756	-0.004380	0.0304	0.008827	0.0712	Yes
	4	0.0416	-0.001682	0.0181	0.005272	0.0399	Yes
	8	0.0245	-0.001725	0.0124	0.003610	0.0228	Yes
	16	0.0143	-0.001217	0.0085	0.002483	0.0131	Yes
	32	0.0084	-0.000938	0.0068	0.001988	0.0075	Yes
	64	0.0038	0.000500	0.0051	0.001496	0.0043	Yes
	128	0.0022	0.000262	0.0048	0.001406	0.0025	Yes
	256	0.0011	0.000290	0.0030	0.000868	0.0014	Yes
	512	-0.0004	0.001220	0.0031	0.000910	0.0008	No
MO-1	$H = 0.75, M = 32, \mu_{on} = 1, \mu_{off} = 127, B = 1024, T = 1024 \times 10^5$						
	1	0.4073	0.006882	0.0964	0.028045	0.4142	Yes
	2	0.2627	0.006923	0.1150	0.033448	0.2696	Yes
	4	0.1829	0.005372	0.1194	0.034709	0.1882	Yes
	8	0.1274	0.005355	0.1194	0.034718	0.1327	Yes
	16	0.0896	0.004195	0.1131	0.032884	0.0938	Yes
	32	0.0646	0.001734	0.1035	0.030090	0.0663	Yes
	64	0.0462	0.000627	0.0895	0.026020	0.0469	Yes
	128	0.0289	0.004205	0.0699	0.020329	0.0331	Yes
	256	0.0149	0.008529	0.0431	0.012531	0.0234	Yes
	512	0.0029	0.013640	0.0135	0.003932	0.0166	No
MO-2	$H = 0.75, M = 16, \mu_{on} = 1, \mu_{off} = 63, B = 1024, T = 1024 \times 10^5$						
	1	0.4410	-0.026778	0.1115	0.032412	0.4142	Yes
	2	0.3031	-0.033452	0.1384	0.040252	0.2696	Yes
	4	0.2239	-0.035641	0.1534	0.044617	0.1882	Yes
	8	0.1700	-0.037291	0.1636	0.047585	0.1327	Yes
	16	0.1338	-0.039984	0.1689	0.049102	0.0938	Yes
	32	0.1053	-0.039047	0.1719	0.049988	0.0663	Yes
	64	0.0853	-0.038446	0.1678	0.048794	0.0469	Yes
	128	0.0663	-0.033201	0.1598	0.046476	0.0331	Yes
	256	0.0513	-0.027851	0.1392	0.040469	0.0234	Yes
	512	0.0321	-0.015555	0.1059	0.030807	0.0166	Yes

Table A.6 Continued: *Complete Autocorrelation Data*

Sim	Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
MO-3	$H = 0.75, M = 8, \mu_{on} = 1, \mu_{off} = 31, B = 1024, T = 1024 \times 10^5$						
	1	0.3884	0.025833	0.0446	0.012979	0.4142	No
	2	0.2372	0.032404	0.0542	0.015752	0.2696	No
	4	0.1552	0.033071	0.0595	0.017313	0.1882	No
	8	0.0981	0.034596	0.0619	0.018013	0.1327	No
	16	0.0621	0.031715	0.0568	0.016522	0.0938	No
	32	0.0359	0.030389	0.0450	0.013086	0.0663	No
	64	0.0179	0.028943	0.0283	0.008234	0.0469	No
	128	0.0068	0.026366	0.0142	0.004116	0.0331	No
	256	0.0024	0.021003	0.0043	0.001261	0.0234	No
	512	0.0006	0.015953	0.0040	0.001162	0.0166	No
MO-4	$H = 0.75, M = 4, \mu_{on} = 1, \mu_{off} = 15, B = 1024, T = 1024 \times 10^5$						
	1	0.3928	0.021383	0.0569	0.016539	0.4142	No
	2	0.2418	0.027875	0.0690	0.020052	0.2696	No
	4	0.1583	0.029903	0.0736	0.021401	0.1882	No
	8	0.1064	0.026318	0.0740	0.021505	0.1327	No
	16	0.0715	0.022231	0.0704	0.020477	0.0938	No
	32	0.0425	0.023773	0.0648	0.018858	0.0663	No
	64	0.0252	0.021705	0.0523	0.015220	0.0469	No
	128	0.0144	0.018722	0.0370	0.010756	0.0331	No
	256	0.0046	0.018867	0.0116	0.003366	0.0234	No
	512	0.0031	0.013511	0.0078	0.002273	0.0166	No
MO-5	$H = 0.75, M = 2, \mu_{on} = 1, \mu_{off} = 7, B = 1024, T = 1024 \times 10^5$						
	1	0.4152	-0.000974	0.0804	0.023387	0.4142	Yes
	2	0.2697	-0.000050	0.0985	0.028639	0.2696	Yes
	4	0.1865	0.001792	0.1093	0.031798	0.1882	Yes
	8	0.1265	0.006188	0.1150	0.033454	0.1327	Yes
	16	0.0854	0.008352	0.1172	0.034075	0.0938	Yes
	32	0.0594	0.006860	0.1178	0.034242	0.0663	Yes
	64	0.0412	0.005683	0.1139	0.033134	0.0469	Yes
	128	0.0308	0.002385	0.1062	0.030889	0.0331	Yes
	256	0.0195	0.003908	0.0962	0.027965	0.0234	Yes
	512	0.0153	0.001224	0.0821	0.023864	0.0166	Yes
MO-6	$H = 0.75, M = 1, \mu_{on} = 1, \mu_{off} = 3, B = 1024, T = 1024 \times 10^5$						
	1	0.4268	-0.012548	0.1273	0.037005	0.4142	Yes
	2	0.2815	-0.011817	0.1589	0.046196	0.2696	Yes
	4	0.1985	-0.010302	0.1767	0.051370	0.1882	Yes
	8	0.1438	-0.011123	0.1870	0.054390	0.1327	Yes
	16	0.1078	-0.014069	0.1927	0.056027	0.0938	Yes
	32	0.0840	-0.017725	0.1956	0.056868	0.0663	Yes
	64	0.0682	-0.021335	0.1958	0.056943	0.0469	Yes
	128	0.0545	-0.021370	0.1945	0.056553	0.0331	Yes
	256	0.0483	-0.024884	0.1914	0.055650	0.0234	Yes
	512	0.0438	-0.027220	0.1832	0.053283	0.0166	Yes

Table A.6 Continued: *Complete Autocorrelation Data*

Sim	Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
MO-7	$H = 0.75, M = 32, \mu_{on} = 1, \mu_{off} = 127, B = 16,384, T = 16,384 \times 10^5$						
	1	0.4191	-0.004885	0.1206	0.035064	0.4142	Yes
	2	0.2765	-0.006863	0.1422	0.041356	0.2696	Yes
	4	0.1927	-0.004482	0.1517	0.044109	0.1882	Yes
	8	0.1374	-0.004687	0.1479	0.043015	0.1327	Yes
	16	0.0992	-0.005417	0.1346	0.039155	0.0938	Yes
	32	0.0661	0.000188	0.1094	0.031808	0.0663	Yes
	64	0.0357	0.011155	0.0817	0.023771	0.0469	Yes
	128	0.0102	0.022917	0.0401	0.011656	0.0331	No
	256	0.0010	0.022449	0.0038	0.001112	0.0234	No
	512	0.0015	0.015084	0.0033	0.000954	0.0166	No
MO-8	$H = 0.75, M = 16, \mu_{on} = 1, \mu_{off} = 63, B = 16,384, T = 16,384 \times 10^5$						
	1	0.3516	0.062587	0.0691	0.020090	0.4142	No
	2	0.1993	0.070323	0.0790	0.022978	0.2696	No
	4	0.1176	0.070639	0.0808	0.023494	0.1882	No
	8	0.0696	0.063087	0.0738	0.021475	0.1327	No
	16	0.0371	0.056627	0.0586	0.017051	0.0938	No
	32	0.0204	0.045912	0.0392	0.011390	0.0663	No
	64	0.0081	0.038802	0.0160	0.004643	0.0469	No
	128	0.0054	0.027721	0.0054	0.001569	0.0331	No
	256	0.0035	0.019971	0.0040	0.001175	0.0234	No
	512	0.0026	0.014019	0.0051	0.001491	0.0166	No
MO-9	$H = 0.75, M = 8, \mu_{on} = 1, \mu_{off} = 31, B = 16,384, T = 16,384 \times 10^5$						
	1	0.4001	0.014159	0.0755	0.021955	0.4142	Yes
	2	0.2481	0.021533	0.0916	0.026623	0.2696	Yes
	4	0.1605	0.027703	0.0955	0.027783	0.1882	Yes
	8	0.1061	0.026662	0.0926	0.026931	0.1327	Yes
	16	0.0683	0.025436	0.0789	0.022950	0.0938	No
	32	0.0369	0.029348	0.0578	0.016821	0.0663	No
	64	0.0144	0.032490	0.0358	0.010423	0.0469	No
	128	0.0067	0.026417	0.0111	0.003240	0.0331	No
	256	0.0035	0.019926	0.0106	0.003069	0.0234	No
	512	0.0024	0.014179	0.0096	0.002792	0.0166	No
MO-10	$H = 0.75, M = 4, \mu_{on} = 1, \mu_{off} = 15, B = 16,384, T = 16,384 \times 10^5$						
	1	0.3524	0.061764	0.0456	0.013267	0.4142	No
	2	0.1999	0.069767	0.0557	0.016188	0.2696	No
	4	0.1172	0.071006	0.0559	0.016266	0.1882	No
	8	0.0650	0.067663	0.0471	0.013690	0.1327	No
	16	0.0356	0.058201	0.0356	0.010341	0.0938	No
	32	0.0164	0.049870	0.0229	0.006660	0.0663	No
	64	0.0078	0.039068	0.0084	0.002451	0.0469	No
	128	0.0049	0.028251	0.0097	0.002818	0.0331	No
	256	0.0026	0.020886	0.0072	0.002108	0.0234	No
	512	0.0007	0.015874	0.0051	0.001494	0.0166	No

Table A.6 Continued: *Complete Autocorrelation Data*

Sim	Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
MO-11	$H = 0.75, M = 2, \mu_{on} = 1, \mu_{off} = 7, B = 16,384, T = 16,384 \times 10^5$						
	1	0.3970	0.017227	0.0946	0.027516	0.4142	Yes
	2	0.2481	0.021501	0.1119	0.032539	0.2696	Yes
	4	0.1672	0.021039	0.1145	0.033304	0.1882	Yes
	8	0.1121	0.020574	0.1116	0.032467	0.1327	Yes
	16	0.0746	0.019212	0.1038	0.030177	0.0938	Yes
	32	0.0544	0.011893	0.0943	0.027431	0.0663	Yes
	64	0.0335	0.013405	0.0780	0.022672	0.0469	Yes
	128	0.0224	-0.022382	0.0700	0.027127	0.0331	Yes
	256	0.0070	-0.006953	0.0169	0.006553	0.0234	No
	512	0.0010	-0.001025	0.0050	0.001948	0.0166	Yes
STR-1	<i>Strictly Alternating</i> $H = 0.75, M = 32, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^5$						
	1	0.4256	-0.011398	0.0143	0.004165	0.4142	No
	2	0.2752	-0.005594	0.0178	0.005185	0.2696	No
	4	0.1893	-0.001094	0.0195	0.005664	0.1882	Yes
	8	0.1314	0.001320	0.0212	0.006178	0.1327	Yes
	16	0.0907	0.003029	0.0217	0.006314	0.0938	Yes
	32	0.0610	0.005269	0.0235	0.006831	0.0663	Yes
	64	0.0419	0.004984	0.0202	0.005863	0.0469	Yes
	128	0.0282	0.004941	0.0198	0.005771	0.0331	Yes
	256	0.0173	0.006173	0.0170	0.004940	0.0234	No
	512	0.0100	0.006570	0.0136	0.003964	0.0166	No
STR-2	<i>Strictly Alternating</i> $H = 0.75, M = 8, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^5$						
	1	0.4246	-0.010366	0.0220	0.006395	0.4142	No
	2	0.2734	-0.003716	0.0268	0.007794	0.2696	Yes
	4	0.1872	0.001028	0.0286	0.008324	0.1882	Yes
	8	0.1275	0.005258	0.0282	0.008199	0.1327	Yes
	16	0.0864	0.007369	0.0291	0.008455	0.0938	Yes
	32	0.0581	0.008185	0.0280	0.008135	0.0663	No
	64	0.0365	0.010346	0.0261	0.007581	0.0469	No
	128	0.0212	0.011941	0.0242	0.007029	0.0331	No
	256	0.0118	0.011596	0.0209	0.006085	0.0234	No
	512	0.0063	0.010248	0.0161	0.004687	0.0166	No
STR-3	<i>Strictly Alternating</i> $H = 0.75, M = 4, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^5$						
	1	0.4213	-0.007131	0.0378	0.010992	0.4142	Yes
	2	0.2682	0.001469	0.0468	0.013597	0.2696	Yes
	4	0.1831	0.005155	0.0519	0.015085	0.1882	Yes
	8	0.1246	0.008100	0.0548	0.015940	0.1327	Yes
	16	0.0839	0.009891	0.0563	0.016370	0.0938	Yes
	32	0.0556	0.010724	0.0560	0.016277	0.0663	Yes
	64	0.0368	0.010035	0.0513	0.014913	0.0469	Yes
	128	0.0250	0.008176	0.0497	0.014449	0.0331	Yes
	256	0.0165	0.006900	0.0453	0.013167	0.0234	Yes
	512	0.0137	0.002871	0.0413	0.012007	0.0166	Yes

Table A.6 Continued: *Complete Autocorrelation Data*

Sim	Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
STR-4	<i>Strictly Alternating</i> $H = 0.75, M = 2, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^5$						
	1	0.4184	-0.004178	0.0293	0.008526	0.4142	Yes
	2	0.2654	0.004261	0.0363	0.010555	0.2696	Yes
	4	0.1791	0.009149	0.0398	0.011565	0.1882	Yes
	8	0.1220	0.010746	0.0426	0.012399	0.1327	Yes
	16	0.0815	0.012230	0.0441	0.012820	0.0938	Yes
	32	0.0512	0.015118	0.0415	0.012070	0.0663	No
	64	0.0304	0.016500	0.0414	0.012033	0.0469	No
	128	0.0191	0.014045	0.0371	0.010800	0.0331	No
	256	0.0086	0.014857	0.0316	0.009203	0.0234	No
	512	0.0053	0.011247	0.0232	0.006759	0.0166	No
STR-5	<i>Strictly Alternating</i> $H = 0.75, M = 32, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^4$						
	1	0.4189	-0.004718	0.0282	0.008207	0.4142	Yes
	2	0.2675	0.002155	0.0352	0.010228	0.2696	Yes
	4	0.1826	0.005628	0.0399	0.011590	0.1882	Yes
	8	0.1234	0.009334	0.0415	0.012081	0.1327	Yes
	16	0.0788	0.014973	0.0451	0.013109	0.0938	No
	32	0.0512	0.015057	0.0415	0.012068	0.0663	No
	64	0.0288	0.018085	0.0402	0.011676	0.0469	No
	128	0.0203	0.012825	0.0320	0.009319	0.0331	No
	256	0.0161	0.007363	0.0267	0.007751	0.0234	Yes
	512	0.0032	0.013387	0.0216	0.006269	0.0166	No
STR-6	<i>Strictly Alternating</i> $H = 0.75, M = 8, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^4$						
	1	0.4273	-0.013065	0.0591	0.017183	0.4142	Yes
	2	0.2740	-0.004400	0.0732	0.021294	0.2696	Yes
	4	0.1898	-0.001534	0.0761	0.022137	0.1882	Yes
	8	0.1265	0.006244	0.0809	0.023528	0.1327	Yes
	16	0.0852	0.008565	0.0802	0.023318	0.0938	Yes
	32	0.0562	0.010122	0.0806	0.023450	0.0663	Yes
	64	0.0376	0.009283	0.0726	0.021099	0.0469	Yes
	128	0.0231	0.010093	0.0614	0.017862	0.0331	Yes
	256	0.0045	0.018933	0.0505	0.014691	0.0234	No
	512	0.0055	0.011101	0.0352	0.010242	0.0166	No
STR-7	<i>Strictly Alternating</i> $H = 0.75, M = 4, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^4$						
	1	0.4058	0.008403	0.0878	0.025526	0.4142	Yes
	2	0.2473	0.022335	0.1094	0.031804	0.2696	Yes
	4	0.1615	0.026737	0.1200	0.034888	0.1882	Yes
	8	0.1049	0.027794	0.1271	0.036953	0.1327	Yes
	16	0.0724	0.021406	0.1287	0.037436	0.0938	Yes
	32	0.0486	0.017694	0.1272	0.036993	0.0663	Yes
	64	0.0304	0.016504	0.1243	0.036153	0.0469	Yes
	128	0.0188	0.014360	0.1178	0.034262	0.0331	Yes
	256	0.0164	0.007045	0.1112	0.032349	0.0234	Yes
	512	0.0225	-0.005923	0.1006	0.029254	0.0166	Yes

Table A.6 Continued: *Complete Autocorrelation Data*

Sim	Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
STR-8	<i>Strictly Alternating</i> $H = 0.75, M = 2, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^4$						
	1	0.4085	0.005712	0.0569	0.016555	0.4142	Yes
	2	0.2538	0.015804	0.0706	0.020527	0.2696	Yes
	4	0.1635	0.024776	0.0723	0.021034	0.1882	No
	8	0.1061	0.026610	0.0732	0.021283	0.1327	No
	16	0.0588	0.034980	0.0646	0.018795	0.0938	No
	32	0.0288	0.037452	0.0511	0.014857	0.0663	No
	64	0.0097	0.037198	0.0269	0.007829	0.0469	No
	128	-0.0015	0.034617	0.0121	0.003508	0.0331	No
	256	-0.0004	0.023849	0.0184	0.005359	0.0234	No
	512	-0.0035	0.020036	0.0124	0.003617	0.0166	No
STR-9	<i>Strictly Alternating</i> $H = 0.85, M = 32, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^5$						
	1	0.6185	0.006037	0.0183	0.005315	0.6245	No
	2	0.4766	0.010873	0.0246	0.007145	0.4875	No
	4	0.3782	0.015200	0.0282	0.008211	0.3934	No
	8	0.3012	0.017790	0.0319	0.009270	0.3190	No
	16	0.2380	0.021026	0.0341	0.009905	0.2590	No
	32	0.1858	0.024559	0.0370	0.010747	0.2104	No
	64	0.1425	0.028416	0.0383	0.011138	0.1709	No
	128	0.1083	0.030519	0.0394	0.011445	0.1388	No
	256	0.0819	0.030832	0.0422	0.012268	0.1127	No
	512	0.0637	0.027843	0.0428	0.012440	0.0916	No
STR-10	<i>Strictly Alternating</i> $H = 0.90, M = 32, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^5$						
	1	0.7141	0.026971	0.0218	0.006348	0.7411	No
	2	0.5902	0.039972	0.0315	0.009149	0.6301	No
	4	0.4964	0.049948	0.0385	0.011201	0.5464	No
	8	0.4171	0.058089	0.0442	0.012844	0.4752	No
	16	0.3487	0.064903	0.0488	0.014189	0.4136	No
	32	0.2871	0.072880	0.0537	0.015628	0.3600	No
	64	0.2364	0.076968	0.0574	0.016699	0.3134	No
	128	0.1906	0.082191	0.0611	0.017771	0.2728	No
	256	0.1528	0.084668	0.0641	0.018653	0.2375	No
	512	0.1203	0.086429	0.0653	0.018999	0.2068	No
STR-11	<i>Strictly Alternating</i> $H = 0.65, M = 32, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^5$						
	1	0.2693	-0.038188	0.0131	0.003822	0.2311	No
	2	0.1428	-0.019466	0.0141	0.004101	0.1233	No
	4	0.0845	-0.010124	0.0130	0.003772	0.0744	No
	8	0.0492	-0.003652	0.0141	0.004091	0.0456	Yes
	16	0.0298	-0.001777	0.0136	0.003954	0.0280	Yes
	32	0.0166	0.000621	0.0143	0.004169	0.0172	Yes
	64	0.0091	0.001554	0.0125	0.003638	0.0106	Yes
	128	0.0048	0.001715	0.0103	0.002996	0.0065	Yes
	256	0.0032	0.000820	0.0084	0.002448	0.0040	Yes
	512	0.0008	0.001650	0.0051	0.001493	0.0025	No

Table A.6 Continued: *Complete Autocorrelation Data*

Sim	Lag	$\hat{\mu}_{r(k)}$	$\Delta\mu$	$\hat{\sigma}_{r(k)}$	$\pm 90\%$ C.I.	Target	Target in C.I.
STR-12	$H = 0.75, M = 32, \mu_{on} = \mu_{off} = 1, B = 1, T = 1024 \times 10^5$						
	1	0.5972	-0.182997	0.0007	0.000201	0.4142	No
	2	0.4723	-0.202604	0.0009	0.000266	0.2696	No
	4	0.3589	-0.170625	0.0011	0.000320	0.1882	No
	8	0.2652	-0.132440	0.0013	0.000375	0.1327	No
	16	0.1923	-0.098505	0.0014	0.000421	0.0938	No
	32	0.1378	-0.071492	0.0015	0.000446	0.0663	No
	64	0.0981	-0.051227	0.0015	0.000445	0.0469	No
	128	0.0695	-0.036396	0.0016	0.000466	0.0331	No
	256	0.0492	-0.025777	0.0017	0.000481	0.0234	No
	512	0.0348	-0.018193	0.0017	0.000495	0.0166	No
D-1	$H = 0.75, M = 32, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^5$ Discr. at 0.01						
	1	0.4148	-0.000549	0.0172	0.005003	0.4142	Yes
	2	0.2671	0.002533	0.0215	0.006243	0.2696	Yes
	4	0.1840	0.004211	0.0234	0.006813	0.1882	Yes
	8	0.1265	0.006178	0.0249	0.007246	0.1327	Yes
	16	0.0877	0.006079	0.0254	0.007397	0.0938	Yes
	32	0.0604	0.005860	0.0271	0.007874	0.0663	Yes
	64	0.0396	0.007268	0.0273	0.007941	0.0469	Yes
	128	0.0257	0.007418	0.0293	0.008514	0.0331	Yes
	256	0.0157	0.007699	0.0269	0.007824	0.0234	Yes
	512	0.0111	0.005493	0.0230	0.006693	0.0166	Yes
D-2	$H = 0.75, M = 32, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^5$ Discr. at 0.1						
	1	0.4176	-0.003370	0.0100	0.002901	0.4142	No
	2	0.2702	-0.000575	0.0135	0.003915	0.2696	Yes
	4	0.1874	0.000810	0.0152	0.004423	0.1882	Yes
	8	0.1309	0.001800	0.0152	0.004430	0.1327	Yes
	16	0.0917	0.002037	0.0155	0.004496	0.0938	Yes
	32	0.0632	0.003073	0.0141	0.004108	0.0663	Yes
	64	0.0420	0.004839	0.0166	0.004835	0.0469	No
	128	0.0284	0.004783	0.0168	0.004885	0.0331	Yes
	256	0.0175	0.005923	0.0156	0.004534	0.0234	No
	512	0.0110	0.005592	0.0144	0.004201	0.0166	No
D-3	$H = 0.75, M = 32, \mu_{on} = \mu_{off} = 1, B = 1024, T = 1024 \times 10^5$ Discr. at 1.0						
	1	0.4254	-0.011149	0.0107	0.003114	0.4142	No
	2	0.2748	-0.005137	0.0145	0.004205	0.2696	No
	4	0.1915	-0.003290	0.0160	0.004661	0.1882	Yes
	8	0.1325	0.000176	0.0174	0.005072	0.1327	Yes
	16	0.0928	0.001005	0.0182	0.005280	0.0938	Yes
	32	0.0648	0.001514	0.0174	0.005064	0.0663	Yes
	64	0.0447	0.002205	0.0194	0.005640	0.0469	Yes
	128	0.0297	0.003468	0.0176	0.005110	0.0331	Yes
	256	0.0192	0.004232	0.0171	0.004965	0.0234	Yes
	512	0.0146	0.001936	0.0142	0.004130	0.0166	Yes

Table A.6 Continued: *Complete Autocorrelation Data*

Appendix B

Source Code Listing

Source Listing B.1: Simulator Source

```
//*****
// File:      sim.c
// Author:    Philip M. Wells
// Description:
//   This program simulates the generation of self-similar network traffic
//   using high-variance, ON/OFF sources. See
//   Wells, P.M., 'Simulation of Self-Similar Network Traffic Using High
//   Variance ON/OFF Sources,' M.S. Thesis, Clemson University, May 2002
//   for more information.
//
// Copyright 2002 Philip M. Wells
//
// This program is free software; you can redistribute it and/or modify
// it under the terms of the GNU General Public License as published by
// the Free Software Foundation; either version 2 of the License, or
// (at your option) any later version.
//
// This program is distributed in the hope that it will be useful,
// but WITHOUT ANY WARRANTY; without even the implied warranty of
// MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
// GNU General Public License for more details.
//
// You should have received a copy of the GNU General Public License
// along with this program; if not, write to the Free Software
// Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
//*****

#include <stdio.h>
#include <stdlib.h>
#include <math.h>
#include <sys/ddi.h>
#include <sys/time.h>

#define ON 1
#define OFF 0

// Runtime parameters
int M; // number of sources
double B; // number of seconds to aggregate << total_time
double H; // Hurst parameter
double mu_on; // mean ON time
double mu_off; // mean OFF time

// Calculated Pareto constants
double alpha_on; // Calculated from H & mean
double alpha_on_neg_inv;
double beta_on;
double alpha_off;
double alpha_off_neg_inv;
```

```

double beta_off;

inline double next_on_pareto() {
    // Using Pareto cumulative dist. of:
    //    $P(X \leq x) = 1 - (\text{beta}/x)^\alpha$ 
    //
    // to generate pseudo random numbers from this dist.
    // take uniform pseudo random numbers of the range [0, 1)
    // and map onto P(x) using the inverse of the cumulative
    // distribution function:
    //    $P\{-1\} = \text{beta} (1 - p) ^ {(-1 / \alpha)}$ 

    return beta_on * pow((1.0 - drand48()), alpha_on_neg_inv);
}

inline double next_off_time() {
    // Return pareto with mean off
    return beta_off * pow((1.0 - drand48()), alpha_off_neg_inv);
}

// Calculate statistics of elements in array W
void stats(double *W, int N)
{
    int i, k;
    int auto_max = 1024; // maximum lag
    double temp;

    double expected_mean, variance, var_normal;
    double norm_factor;
    double *r; // autocorrelations array

    expected_mean = mu_on / (mu_on + mu_off) * M * B;

    // calc variance and normalized variance
    norm_factor = 1 / (pow(B, 2 * H) * M);
    variance = 0;
    for (i = 0; i < N; i++) {
        temp = (W[i] - expected_mean);
        variance += temp * temp;
    }
    variance = variance / (double) N;
    var_normal = variance * norm_factor;

    fprintf(stderr, "%lf ", var_normal);

    // limit the lag if there are not enough data points
    if (N/2 < auto_max)
        auto_max = N/2;
    r = malloc(sizeof(double) * auto_max);

    // Calculate auto correlation values
    r[0] = 0;
    for (k = 1; k < auto_max; k++) {
        for (i = 0; i < N-k; i++) {
            r[k] += (W[i] - expected_mean) * (W[i+k] - expected_mean);
        }
        r[k] = r[k] / (double) (N - k);
        r[k] = r[k] / variance;
    }

    for (i = 0; i < auto_max; i++)
        printf("%lf\n", r[i]);
}

```

```

}

int main (int argc, char *argv[])
{
    double T; // total simulation time
    int N; // number of time steps
#ifdef DISCRETE
    double discr_size;
#endif

    double *cur_source_times;
    unsigned char *source_state;
    double *pack_arrival_count; // Aggregated packet arrival count array

    double on_time_this_block;
    double cur_sim_time;
    double next_sim_time;
    double new_on_time;

    struct timeval cur_timeofday;

    int i, proc, cur_block;

    // Extract args
#ifdef DISCRETE
    if (argc == 8) {
        H = atof(argv[1]);
        mu_on = atof(argv[2]);
        mu_off = atof(argv[3]);
        M = atoi(argv[4]);
        B = atof(argv[5]);
        T = atof(argv[6]);
        discr_size = atof(argv[7]);
    } else {
        fprintf(stderr,
            "Usage: sim <H> <mean_on> <mean_off> <M> <B> <T> <discr_time>\n");
        exit(1);
    }
#else
    if (argc == 7) {
        H = atof(argv[1]);
        mu_on = atof(argv[2]);
        mu_off = atof(argv[3]);
        M = atoi(argv[4]);
        B = atof(argv[5]);
        T = atof(argv[6]);
    } else {
        fprintf(stderr,
            "Usage: sim <H> <mean_on> <mean_off> <M> <B> <T>\n");
        exit(1);
    }
#endif

    // Seed random number generator
    gettimeofday(&cur_timeofday, NULL);
    srand48(cur_timeofday.tv_usec);

    // Set pareto constants
    alpha_on = 3.0 - 2.0 * H;
    beta_on = (mu_on * (alpha_on - 1)) / alpha_on;
    alpha_on.neg_inv = -1.0 / alpha_on;

    alpha_off = 3.0 - 2.0 * H;

```

```

    beta_off = (mu_off * (alpha_off - 1)) / alpha_off;
    alpha_off_neg_inv = -1.0 / alpha_off;

    // Create cur_source_times array
    N = T / B;
    pack_arrival_count = malloc(N * sizeof(double));
    cur_source_times = malloc(M * sizeof(double));
    source_state = malloc(M * sizeof(unsigned char));

    // Start sources at different times
    for (i = 0; i < M; i++) {

        // Give each source a random, uniformly dist. starting time before time
        // zero - this range is just a guess...
#ifdef DISCRETE
        cur_source_times[i] =
            (double)((int)(-10 * (mu_on+mu_off) * M * drand48() / discr_size)) *
            discr_size;
#else
        cur_source_times[i] = -10 * (mu_on+mu_off) * M *
            drand48();
#endif
        source_state[i] = OFF;

        while (cur_source_times[i] <= 0) {

#ifdef STRICTALTERNATE
            // Pick random current state
            source_state[i] = (drand48() > 0.5) ? ON : OFF;
#endif

            // Do on
            if (source_state[i] == OFF) {
#ifdef DISCRETE
                cur_source_times[i] +=
                    (double)((int)(next_on_pareto() / discr_size)) * discr_size;
#else
                cur_source_times[i] += next_on_pareto();
#endif
                source_state[i] = ON;
            }
            // Do off
            else {
#ifdef DISCRETE
                cur_source_times[i] +=
                    (double)((int)(next_off_time() / discr_size)) * discr_size;
#else
                cur_source_times[i] += next_off_time();
#endif
                source_state[i] = OFF;
            }
        }
    }

    cur_sim_time = 0; // Beginning of the current time block
    for (cur_block = 0; cur_block < N; cur_block++)
    {
        on_time_this_block = 0;
        next_sim_time = cur_sim_time + B;

        for (proc = 0; proc < M; proc++)
        {

```



```

// If source is currently on, add its time since the beginning of
// the time block
if (source_state[proc] == ON)
{
    new_on_time = cur_source_times[proc] - cur_sim_time;
    if (new_on_time > B)
        on_time_this_block += B;
    else
    {
#ifdef DISCRETE
        on_time_this_block +=
            (double)((int)(new_on_time / discr_size)) * discr_size;
#else
        on_time_this_block += new_on_time;
#endif
    }
}

while (cur_source_times[proc] <= next_sim_time)
{
#ifdef STRICTALTERNATE
    // Pick random current state
    source_state[proc] = (drand48() > 0.5) ? ON : OFF;
#endif

    // Do on
    if (source_state[proc] == OFF)
    {
        // Add this on time -- until end of block -- to
        // on_time_this_block
        new_on_time = next_on_pareto();
        if (new_on_time + cur_source_times[proc] <
            next_sim_time)
        {
#ifdef DISCRETE
            on_time_this_block +=
                (double)((int)(new_on_time / discr_size)) *
                discr_size;
#else
            on_time_this_block += new_on_time;
#endif
        }
        else
            on_time_this_block += next_sim_time -
                cur_source_times[proc];

#ifdef DISCRETE
        cur_source_times[proc] +=
            (double)((int)(new_on_time / discr_size)) * discr_size;
#else
        cur_source_times[proc] += new_on_time;
#endif
        source_state[proc] = ON;
    }
    // Do off
    else {
#ifdef DISCRETE
        cur_source_times[proc] +=
            (double)((int)(next_off_time() / discr_size)) *
            discr_size;
#else
        cur_source_times[proc] += next_off_time();
#endif
        source_state[proc] = OFF;
    }
}

```

```
    }  
}  
  
pack_arrival_count[cur_block] = on_time_this_block;  
  
cur_sim_time = next_sim_time;  
  
}  
  
stats(pack_arrival_count, N);  
}
```

Bibliography

- BERAN, J., SHERMAN, R., TAQQU, M., AND WILLINGER, W. 1995. Long-range dependence in variable-bit-rate video traffic. *IEEE Transactions on Communications* 43, 1566–1579.
- CAO, J., CLEVELAND, W. S., LIN, D., AND SUN, D. X. 1999. Internet traffic tends to poisson and independent as the load increases. Tech. rep., Bell Labs.
- CHE, H. AND LI, S. 1997. Fast algorithms for measurement-based traffic modeling. In *INFOCOM (1)*. 177–186.
- CHRISTIANSEN, M., JAFFAY, K., OTT, D., AND SMITH, F. D. 2000. Tuning RED for web traffic. In *SIGCOMM*. 139–150.
- CROVELLA, M. E. AND BESTAVROS, A. 1997. Self-similarity in World Wide Web traffic: evidence and possible causes. *IEEE/ACM Transactions on Networking* 5, 6, 835–846.
- ERRAMILI, A., NARAYAN, O., AND WILLINGER, W. 1996. Experimental queueing analysis with long-range dependent packet traffic. *IEEE/ACM Transactions on Networking* 4, 2, 209–223.
- GEIST, R. M. AND WESTALL, J. M. 2000. Practical aspects of simulating systems having arrival processes with long-range dependence. In *Proceedings of the 2000 Winter Simulation Conference*. 666–674.
- JAIN, R. AND ROUTHIER, S. A. 1986. Packet trains — measurements and a new model for computer network traffic. *IEEE Journal on Selected Areas in Communications* 4, 6 (September), 986–995.
- KLIVANSKY, S. M., MUKHERJEE, A., AND SONG, C. 1994. On long-range dependence in NSFNET traffic. Tech. Rep. GIT-CC-94-61.
- LELAND, W. E., TAQQU, M. S., WILLINGER, W., AND WILSON, D. V. 1994. On the self-similar nature of Ethernet traffic (extended version). In *IEEE/ACM Transactions on Networking*. Vol. 2. 1–15.
- MANDLEBROT, B. B. 1969. Long-run linearity, locally gaussian processes, h-spectra and infinite variance. *International Economic Review* 10, 82–113.
- NIKOLAIDIS, I., COOPER, C. A., PERUMALLA, K. S., AND FUJIMOTO, R. 1997. Time-parallel generation of self-similar ATM traffic. In *Winter Simulation Conference*. 1071–1078.
- PARK, K., KIM, G., AND CROVELLA, M. 1997. On the effect and control of self-similar network traffic: A simulation perspective. In *Proceedings of the SPIE International Conference on Performance and Control of Network Systems*.
- PAXSON, V. 1995. Fast approximation of self similar network traffic. *ACM SIGCOMM Computer Communications Review* 27, 5, 5–18.
- PAXSON, V. AND FLOYD, S. 1995. Wide area traffic: the failure of Poisson modeling. *IEEE/ACM Transactions on Networking* 3, 3, 226–244.
- POPESCU, A. 1999. Traffic self-similarity. In *compendium University of Karlskrona/Ronneby*.
- TAQQU, M. S., WILLINGER, W., AND SHERMAN, R. 1997. Proof of a fundamental result in self-similar traffic modeling. *ACMCCR: Computer Communication Review* 27, 5–23.

- WILLINGER, W., TAQQU, M. S., SHERMAN, R., AND WILSON, D. V. 1995. Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level. *Computer Communications Review* 25, 100–113.
- WILLINGER, W., TAQQU, M. S., SHERMAN, R., AND WILSON, D. V. 1997. Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level. *IEEE/ACM Transactions on Networking* 5, 1, 71–86.