Proceedings of the 2002 Winter Simulation Conference
E. Yücesan, C.-H. Chen, J. L. Snowdon, and J. M. Charnes, eds.

A HIGHLY EFFICIENT M/G/∞ MODEL FOR GENERATING SELF-SIMILAR TRACES

María Estrella Sousa-Vieira
Andrés Suárez-González
Cándido López-García
Manuel Fernández-Veiga
José C. López-Ardao

ETSE de Telecomunicación, Universidade de Vigo
36200 Vigo, SPAIN

ABSTRACT

Several traffic measurement reports have convincingly shown the presence of self-similarity in modern networks, inducing as a result a revolution in the stochastic modeling of traffic. The use of self-similar processes in performance analysis has opened new problems and research issues in simulation studies, where the efficient generation of synthetic sample paths with self-similar properties is one of the fundamental concerns. In this paper, we present an M/G/∞ generator of self-similar traces, based on a highly efficient simulation model using the decomposition property of Poisson processes.

1 INTRODUCTION

Several traffic measurement studies (Leland et al. 1994, Garrett and Willinger 1994) have demonstrated the existence of statistical self-similarity in modern networks, along with a closely related property called Long-Range Dependence (LRD), that involves non negligible correlations over arbitrarily large time scales. These findings have contributed to a very important revolution in the stochastic modeling of traffic, since the presence of LRD may have a drastic impact on the performance metrics (Likhanov, Tsybakov and Georganas 1995, Erramilli et al. 1996), and the validity of traditional processes, like Markovian or Autoregressive, is in doubt because modeling LRD through these processes requires many parameters, whose interpretation becomes difficult.

Because of this, the use of self-similar processes for network traffic modeling purposes is essential, due to their capability to exhibit LRD over all time scales by making use of few parameters (parsimonious modeling).

The application of self-similar processes in network simulation studies has opened a wide range of research topics dealing with new problems. One of the most important issues is the synthetic generation of sample paths of LRD processes, since real traces collected by measurements are of limited length and lack the necessary diversity required to make flexible enough simulation studies.

A very interesting self-similar process is the occupancy process of an M/G/∞ queueing model, referred to as M/G/∞ process. It belongs to the class of LRD processes when G, the distribution of the service time, is heavy-tailed of infinite variance.

Apart from its use in analytical studies, the M/G/∞ process has several important advantages for simulation studies, such as the possibility of on-line generation. Furthermore, there exists a trivial method of producing exact sample paths of the process with complexity $\mathcal{O}(n)$, being n the length of the sample path: it suffices to simulate the M/G/∞ queue, sampling the occupancy of the system at integer instants.

Varying the service time distribution, G, many forms of time dependence can be obtained, which makes this process a good candidate for modeling many types of correlated traffic, such as video traffic (Krunz and Makowski 1998). In Suárez et al. (2002) the authors present a discrete random variable whose distribution (S distribution) is heavy-tailed with two parameters, a feature that enables the modeling of both short-term and long-term correlation behavior of the resulting M/S/∞ process.

Despite its high flexibility, the marginal distribution of the M/G/∞ process is Poisson, which is not adequate to model the empirical marginal distribution of some real video sequences. So, we need to transform the Poisson marginal distribution of the M/G/∞ process into a more appropriate heavy-tailed form. However, small values of the arrival rate λ of the Poisson input process are inappropriate for the transformation process (Poon and Lo 2001) and, on the other hand, the complexity of the generator is an increasing function of λ.

In order to get a more efficient simulation model, in this paper we propose the use of the decomposition property of Poisson processes. With the proposed method, we obtain a flexible and efficient M/G/∞ generator able to deal with a wide range of input parameters.

The remainder of the paper is organized as follows. In Section 2 we review the main concepts related to LRD and statistical self-similarity and present the M/G/∞ process. In Section 3 we present the method that we propose in order to improve the efficiency of the generator of samples of the M/G/∞ process, and evaluate the improved simulation model applied to the sample generation of the M/S/∞ process (the S distribution is described in Appendix A). Finally, Section 4 summarizes the conclusions of the paper.

2 LRD AND M/G/∞ PROCESS

Let $X = \{X_k; k = 1, 2, \dots\}$ be a stationary stochastic process with finite variance and autocorrelation function $\mathrm{r}_k \triangleq \mathrm{Cov}\left[X_i, X_{i+k}\right] / \mathrm{Var}\left[X\right]$ and let $X^{(m)}$ be the corresponding aggregated process (with aggregation level m), obtained by averaging the original sequence X over non-overlapping blocks of size m, $X^{(m)} = \left\{\overline{X_i}[m]; i = 1, 2, \dots\right\}$, where:

$$\overline{X_i}[m] = \frac{1}{m} \cdot \sum_{j=(i-1)\cdot m+1}^{i\cdot m} X_j .$$

It is said that X exhibits LRD when its autocorrelation function is not summable, i.e., $\sum_{k=1}^{\infty} \mathrm{r}_k = \infty$, like in those processes whose autocorrelation function decays hyperbolically:

$$\exists \beta \in (0, 1) \left| \lim_{k\to\infty} \frac{\mathrm{r}_k}{k^{-\beta}} = \mathsf{c}_\mathrm{r} \in (0, \infty) \right. . \tag{1}$$

The process X is called exactly second-order self-similar, with self-similarity parameter H (Hurst 1951), if the aggregated process $X^{(m)}$ scaled by $m^{1-\mathsf{H}}$ has the same variance and autocorrelation as X for all m, that is, if the aggregated processes possess the same nondegenerate correlation structure as the original stochastic process. The autocorrelation function of both X and $X^{(m)}$ is:

$$\mathrm{r}_k = \mathrm{g}_k^{\mathsf{H}} \triangleq \frac{1}{2} \cdot \left[(k+1)^{2\mathsf{H}} - 2k^{2\mathsf{H}} + (k-1)^{2\mathsf{H}}\right] \quad \forall k \geq 1 \tag{2}$$

where for $1/2 < \mathsf{H} < 1$ (Cox 1984):

$$\lim_{k\to\infty} \frac{\mathrm{g}_k^{\mathsf{H}}}{k^{2\mathsf{H}-2}} = \mathsf{H} \cdot (2\mathsf{H} - 1),$$

that is, it decays hyperbolically as in (1), and so the process exhibits LRD.

If expression (2) is satisfied asymptotically by the autocorrelation function of the aggregated process, $\mathrm{r}_k^{(m)}$, then the process is called asymptotically second-order self-similar:

$$\lim_{m\to\infty} \mathrm{r}_k^{(m)} = \mathrm{g}_k^{\mathsf{H}} \quad \forall k \geq 1.$$

It has been shown that a covariance stationary process whose autocorrelation function decays hyperbolically is asymptotically second-order self-similar (Tsybakov and Georganas 1997).

A quite interesting self-similar process is the occupancy process of an M/G/∞ queueing model. In such a queueing model, customers arrive according to a Poisson process with rate λ to a pool of infinitely many servers, and their service times constitute a sequence of continuous i.i.d. random variables distributed as the random variable S of finite mean value.

Cox and Isham (1980) showed that the number of customers, or busy servers, in the system at any instant t, $\{X(t); t \in \Re\}$, has a Poisson marginal distribution. If the mean value of S is finite, it can be demonstrated that the occupancy process exhibits LRD iff its variance is infinite, as it may happen in heavy-tailed service distributions.

We are interested on the discrete-time version of $\left\{X(t); t \in \Re^+\right\}$, that is: $X \triangleq \left\{X_i \triangleq X(i); i = 1, 2, \dots\right\}$, stochastic process referred to as the M/G/∞ process.

The most natural approach to generate an M/G/∞ process is to use a discrete-time model, since its simulation will be more efficient (Suárez et al. 2002).

2.1 Discrete-Time Model

Let $A = \{A_n; n = 1, 2, \dots\}$ be a renewal stochastic process, where A_n is a Poisson random variable with mean value λ and represents the number of arrivals at instant n; let $\left\{\left\{S_{n,i}; i = 1, \dots, A_n\right\}; n = 1, 2, \dots\right\}$ be a renewal stochastic process where $S_{n,i}$ is distributed as a positive-valued discrete random variable S with finite mean value $\mathrm{E}[S]$, and corresponds to the service time of the i-th arrival at instant n.

If the following conditions hold:

- the initial number of users X_0 is a Poisson random variable of mean value $\lambda \cdot \mathrm{E}[S]$;
- the service times of these X_0 initial users $\left\{\widehat{S}_j;\ j = 1, \dots, X_0\right\}$ are mutually independent and have the same distribution as the residual life of S, $\widehat{S}$:

$$\Pr\left[\widehat{S} = k\right] = \frac{\Pr\left[S \geq k\right]}{\mathrm{E}[S]},$$

then the stochastic process $X = \{X_n; n = 1, 2, \dots\}$ is strict-sense stationary and ergodic, and enjoys equivalent properties to those of the original continuous-time M/G/∞ process (Parulekar and Makowski 1996):

1. the process X has a Poisson marginal distribution and mean value:
$$\mu \triangleq \mathrm{E}[X] = \lambda \cdot \mathrm{E}[S], \tag{3}$$
2. its autocorrelation function is given by:
$$\begin{aligned} \mathrm{r}_k &= \frac{1}{\mathrm{E}[S]} \cdot \sum_{i=k}^{\infty} \Pr[S > i] \\ &= 1 - \frac{\sum_{i=0}^{k-1} \Pr[S > i]}{\mathrm{E}[S]} \quad \forall k = 1, 2, \ldots \end{aligned} \tag{4}$$
3. it exhibits LRD $\iff \mathrm{E}\left[\widehat{S}\right] = \infty$, since:
$$\sum_{k=0}^{\infty} \mathrm{r}_k = \mathrm{E}\left[\widehat{S}\right].$$

Suárez et al. (2002) propose to use the S distribution for the service time, since its mean value and the cdf of its residual life have explicit expressions, given in Appendix A. Moreover, it is a heavy-tailed distribution with two parameters, α and m, a feature that enables the modeling of both short-term and long-term correlation behavior of the occupancy process. Specifically, the autocorrelation function of the resulting M/S/∞ process is:

$$\mathrm{r}_k = \begin{cases} 1 - \frac{\alpha - 1}{\mathsf{m}\alpha} \cdot k & \forall k \in (0, \mathsf{m}] \\ \frac{1}{\alpha} \cdot \left(\frac{\mathsf{m}}{k}\right)^{\alpha - 1} & \forall k \geq \mathsf{m}. \end{cases}$$

Given the three desired parameters of the process X (mean value μ, Hurst parameter H and one-lag autocorrelation coefficient r_1) the parameters of the M/S/∞ model can be computed as follows:

$$\alpha = 3 - 2\mathsf{H} \tag{5}$$

$$\mathsf{m} = \begin{cases} (\alpha \mathrm{r}_1)^{\frac{1}{\alpha - 1}} & \forall \mathrm{r}_1 \in \left(0, \frac{1}{\alpha}\right] \\ \frac{1 - \frac{1}{\alpha}}{1 - \mathrm{r}_1} & \forall \mathrm{r}_1 \in \left[\frac{1}{\alpha}, 1\right) \end{cases}$$

$$\lambda = \begin{cases} \mu \cdot \frac{\alpha \mathsf{m} - \mathsf{m}^{\alpha}}{\mathsf{m}\alpha} & \forall \mathsf{m} \in (0, 1] \\ \mu \cdot \frac{\alpha - 1}{\mathsf{m}\alpha} & \forall \mathsf{m} \geq 1. \end{cases}$$

3 AN EFFICIENT M/G/∞ GENERATOR

The tail of the marginal distribution plays an important role in performance evaluation (Grossglauser and Bolot 1996). The M/G/∞ process has Poisson marginal distribution, whose tail drops faster than that of the empirical marginal distribution of some real sequences. Therefore, we need to transform the Poisson marginal distribution of the M/G/∞ process into a more appropriate one, but this introduces an efficiency problem. On the one hand, small values of the arrival rate λ of the Poisson input process are inappropriate for the transformation operation but, on the other hand, the complexity of the generator is a linear increasing function of λ.

In order to improve the efficiency for large mean values of the M/G/∞ process, which in view of equation (3) implies also large values of λ, we propose to use the decomposition property of Poisson processes.

3.1 Description of the Proposed Method

As we have seen in Section 2, when we use a discrete-time simulation model of the M/G/∞ system, every sample value X_n requires the generation of:

- one sample of the Poisson random variable A_n, with mean value λ.
- A_n samples of the random variable S.

We denote by N the mean number of random values that have to be generated for each sample value of the occupancy process. In this case $\mathsf{N} = \lambda + 1$. For large values of λ, the computational time can be very high.

In order to improve the efficiency, we divide the arrivals at each instant n into $\mathsf{K} + 1$ groups, according to the values of their service times. The mean number of arrivals at each group is $\lambda \cdot \Pr[S = i]$; $i = 1, 2, \ldots, \mathsf{K}$ for the K first groups and $\lambda \cdot \Pr\left[S > \mathsf{K}\right]$ for the last group.

As we can see in Figure 1, using the decomposition property of Poisson processes we divide the original arrivals input process into $\mathsf{K} + 1$ Poisson processes. Hence, it will produce a statistically identical process if, per sample value X_n, we only generate one sample of the arrivals random variable for each one of the K first groups, since the service times are directly $i = 1, 2, \ldots, \mathsf{K}$.

With this method, each sample value X_n requires the generation of:

- one sample per Poisson random variable $A_{n,i}$; $i = 1, 2, \ldots, \mathsf{K}$, with mean value $\lambda \cdot \Pr[S = i]$ for $i = 1, 2, \ldots, \mathsf{K}$,
- one sample of the Poisson random variable $A_{n,\mathsf{K}+1}$, with mean value $\lambda \cdot \Pr\left[S > \mathsf{K}\right]$,

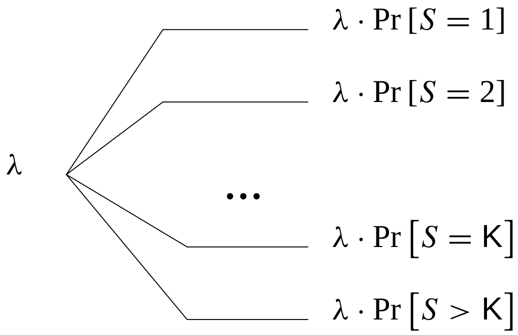

Figure 1: The Arrivals Input Process

- $A_{n,\mathsf{K}+1}$ samples of the random variable $S|_{S>\mathsf{K}}$, that is, of the part of S greater than K.

The mean number of random values that have to be generated for each sample value of the occupancy process is $\mathsf{N} = \mathsf{K} + 1 + \lambda \cdot \Pr[S > \mathsf{K}]$.

Since our aim is to minimize this quantity, N, we choose K as the highest value such that $\lambda \cdot \Pr[S = \mathsf{K}] > 1$. This will ensure the minimum N whenever $\Pr[S = i]$ is a monotonically decreasing function of i.

3.2 Performance for the M/S/∞ Process

We evaluate this improved simulation model with the M/S/∞ process. In the remaining of this section, we check the effect of the input parameters (mainly the mean value μ, but also the Hurst parameter H and the one-lag autocorrelation coefficient r_1) on both the threshold K and the mean number of random values per sample value X_n, N. In the next section we will measure the real performance of our implementation of the proposed simulation model (briefly commented on in Appendix B).

In the following, we vary the mean value of the M/S/∞ process in powers of two, $\mu = 2^m$ (the axis will be in logarithmic scale), and fix one of the two remaining input parameters to a moderate value ($r_1 = 0.5$, $\mathsf{H} = 0.7$), while using a set of values for the other one ($\mathsf{H} \in \{0.6, 0.7, 0.8, 0.9\}$, $r_1 \in \{0.1, 0.5, 0.9\}$): the higher (lower) values for r_1 and H are meant to be representative of strong (weak) correlation and of strong (weak) LRD behavior.

We have used a logarithmic scale for those figures showing the effect on N since it quickly takes large values as μ increases.

First, we show the effect of the mean value of the M/S/∞ process and the Hurst parameter on both K and N in Figures 2 and 3 respectively. With respect to μ, we observe how K takes moderate values even for quite high values of μ, while N is roughly two orders of magnitude lower than μ in the studied interval. On the other hand, we can see that K and N are neither greatly influenced by H, nor they are increasing functions of H for every μ, although asymptotically this seems to be the case. This fact may seem counterintuitive, since it may appear that a higher K should give rise to a lower N. Nevertheless, it is simply the result of the higher dispersion of the service time distribution for lower α (higher H) values while not varying its mean value $\mathrm{E}[S]$: note from equation (4) that the same r_1 implies the same $\mathrm{E}[S]$.

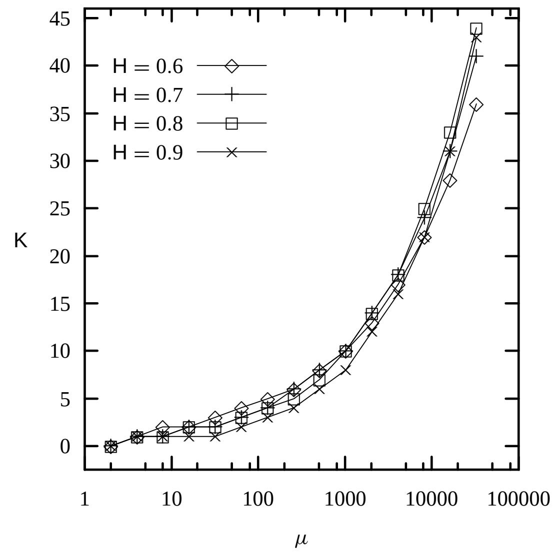

Figure 2: K ($r_1 = 0.5$)

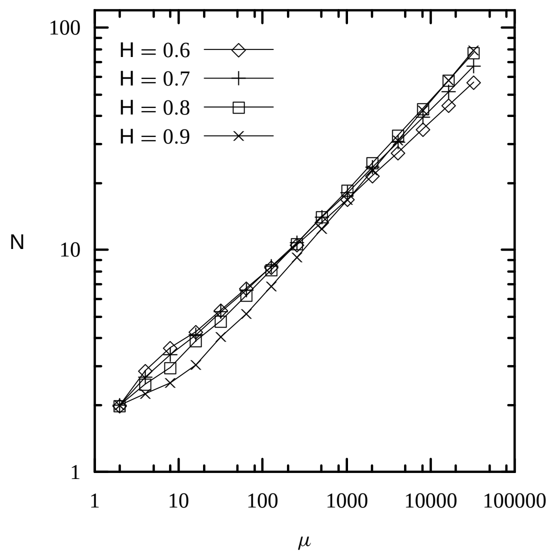

Figure 3: N ($r_1 = 0.5$)

As expected, Figures 4 and 5 reproduce the same behavior of K and N with respect to μ. On the other hand, the effect of the parameter r_1 on K and N is clearly stronger than that of H, with both K and N being increasing functions of r_1 for large values of μ.

This behavior is a consequence of a constant dispersion of S —same H implies same α, from equation (5)— and:

- $\mathrm{E}[S]$ increases with r_1, from expression (4),
- λ decreases as $\mathrm{E}[S]$ increases, from equation (3).

In Figure 4, we observe how K is mainly driven by the behavior of $\mathrm{E}[S]$ instead of that of λ.

Figure 5 depicts the shape in which N increases as r_1 does, for large enough μ. This is mainly due to the increasing K, since both K and the decreasing λ are capable to compensate the increasing $\mathrm{E}[S]$ in $\lambda \cdot \Pr\left[S > \mathsf{K}\right]$.

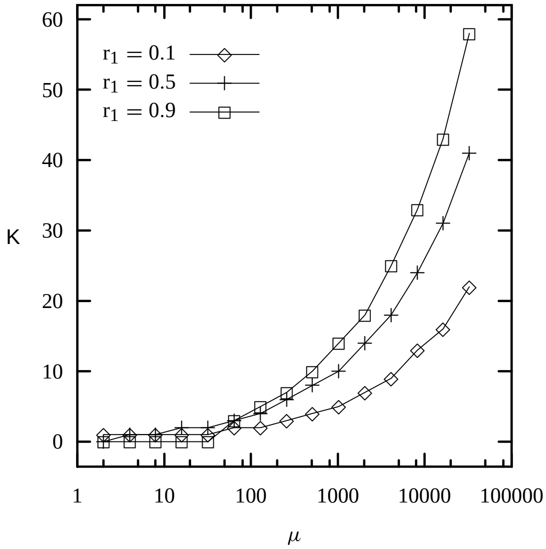

Figure 4: K ($\mathsf{H} = 0.7$)

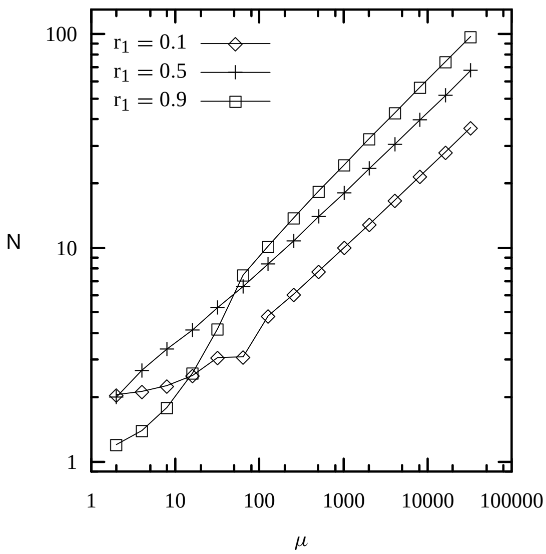

Figure 5: N ($\mathsf{H} = 0.7$)

3.3 Empirical Case Study

In this section we measure the performance of the improved M/S/∞ simulation model, both absolute and relative to that of the previous simulation model (Suárez et al. 2002):

$$\mathsf{R} = \frac{\text{running time of previous version}}{\text{running time of improved version}}.$$

We measured the running times in seconds (using the Unix command time) for the generation of sample paths of $n = 2 \cdot 10^7$ values of the M/S/∞ process in an Athlon XP 1600+ @1.4GHz.

In Figure 6 we observe how the efficiency of the generator based on the improved version is practically equal to that of the previous one for small values of μ, but significantly better as μ increases. This is mainly due to the reduction in the number of samples that we have to generate for each sample value of the occupancy process: $\mathsf{N} = \lambda + 1$ with the previous simulation model and $\mathsf{N} = \mathsf{K} + 1 + \lambda \cdot \Pr\left[S > \mathsf{K}\right]$ with the improved version.

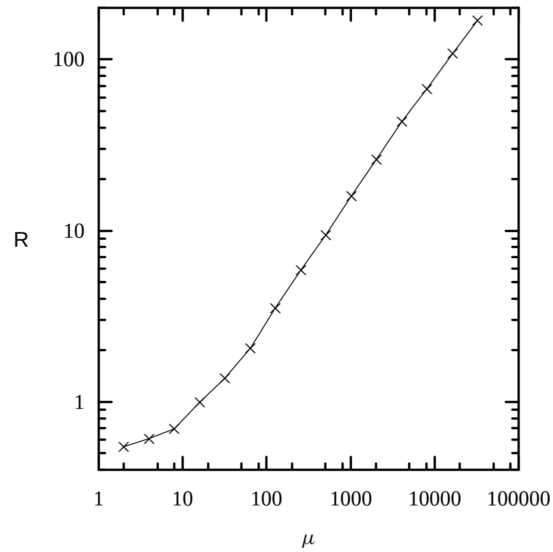

Figure 6: CPU Time Improvement

Table 1 shows numerically the improvement factor, R, for different values of μ, r_1 and H.

Table 1: CPU Time Improvement

M/S/∞ parameters		$\mathsf{H} = 0.6$	$\mathsf{H} = 0.9$
$r_1 = 0.1$	$\mu = 8192$	189.74	185.71
	$\mu = 16384$	314.89	279.63
	$\mu = 32768$	501.69	402.63
$r_1 = 0.9$	$\mu = 8192$	9.375	7.93
	$\mu = 16384$	15.08	8.58
	$\mu = 32768$	21.55	6.55

Both the effect of μ, and that of H or r_1 on the running time in seconds of the improved M/S/∞ model are shown in Figures 7 and 8 respectively. Comparing them to Figures 3 and 5, we see how the running time behaves as a function of the mean number of random values per sample value X_i, N, as expected.

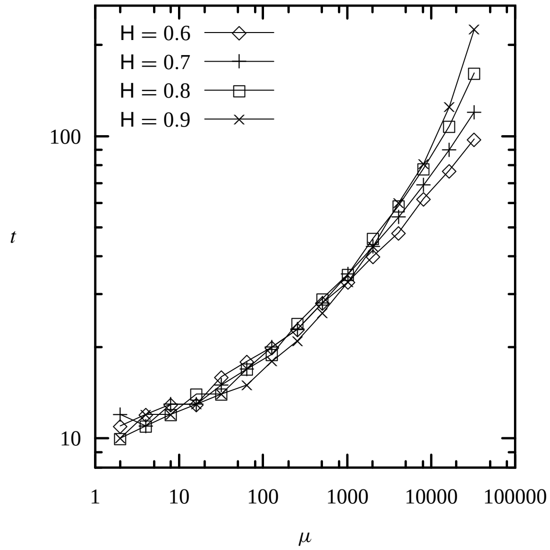

Figure 7: CPU Time ($r_1 = 0.5$)

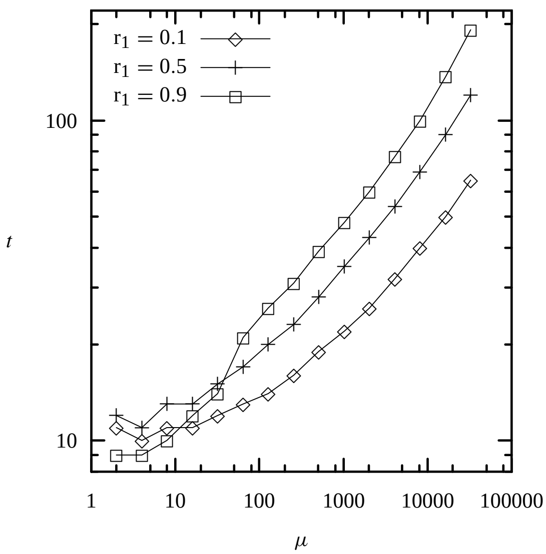

Figure 8: CPU Time ($\mathsf{H} = 0.7$)

We can see that the time is an increasing function of μ (and so of λ) but the complexity is sub $\mathcal{O}(\mu \cdot n)$, although it also depends on the other two parameters, H and r_1.

4 CONCLUSIONS

We have presented in this paper a highly efficient simulation model to obtain an M/G/∞ generator of self-similar traces flexible enough to deal with a wide range of input parameters.

The model is based on the use of the decomposition property of Poisson processes, in order to minimize the mean number of random values to be generated per sample value of the occupancy process.

We have checked both analytically and experimentally the efficiency of the simulation model applied to the generation of samples of the M/S/∞ process.

We are currently investigating how to apply the marginal distribution change technique proposed by Crouse and Baraniuk (1999) for Gaussian marginal distribution to the Poisson one, in order to not mess up too much the autocorrelation structure of the target process.

ACKNOWLEDGMENTS

This work was partially supported by the project TIC2000-1126 of the *Plan Nacional de Investigación Científica, Desarrollo e Innovación Tecnológica*.

APPENDIX A: S DISTRIBUTION

Considering two separate intervals for the parameter m, the distribution of the S discrete-time random variable $\Pr[S = k]$ (Suárez et al. 2002) is, for $\mathsf{m} \le 1$:

- $1 + \frac{\mathsf{m}^\alpha}{\alpha\mathsf{m} - \mathsf{m}^\alpha} \cdot \left[(k+1)^{1-\alpha} - k^{1-\alpha}\right]$ $k = 1$
- $\frac{\mathsf{m}^\alpha}{\alpha\mathsf{m} - \mathsf{m}^\alpha} \cdot \left[(k+1)^{1-\alpha} - 2 \cdot k^{1-\alpha} + (k-1)^{1-\alpha}\right]$ $\forall k > 1$,

and for $\mathsf{m} > 1$:

- $1 + k - \mathsf{m} + \frac{\mathsf{m}^\alpha}{\alpha - 1} \cdot \left[(k+1)^{1-\alpha} - \mathsf{m}^{1-\alpha}\right]$ $k = \lfloor \mathsf{m} \rfloor$
- $1 + \mathsf{m} - k + \frac{\mathsf{m}^\alpha}{\alpha - 1} \cdot \left[(k+1)^{1-\alpha} - 2k^{1-\alpha} + \mathsf{m}^{1-\alpha}\right]$ $k = \lceil \mathsf{m} \rceil$
- $\frac{\mathsf{m}^\alpha}{\alpha - 1} \cdot \left[(k+1)^{1-\alpha} - 2k^{1-\alpha} + (k-1)^{1-\alpha}\right]$ $\forall k > \lceil \mathsf{m} \rceil$.

Its mean value is:

$$\mathrm{E}[S] = \begin{cases} \dfrac{\alpha\mathsf{m}}{\alpha\mathsf{m} - \mathsf{m}^\alpha} & \forall \mathsf{m} \in (0, 1] \\ \dfrac{\alpha\mathsf{m}}{\alpha - 1} & \forall \mathsf{m} \ge 1. \end{cases}$$

As it would be expected, this proposed distribution satisfies the heavy-tailed condition (Leland et al. 1994):

$$\lim_{k\to\infty} \frac{\Pr[S>k]}{k^{-\alpha}} = \mathsf{c}_{\mathrm{h}} \in (0,\infty),$$

with parameter:

$$\mathsf{c}_{\mathrm{h}} = \begin{cases} \mathsf{m}^{\alpha} \cdot \dfrac{\alpha-1}{\alpha\mathsf{m}-\mathsf{m}^{\alpha}} & \forall \mathsf{m} \leq 1 \\ \mathsf{m}^{\alpha} & \forall \mathsf{m} \geq 1, \end{cases}$$

The cdf of the residual life of S is:

$$\Pr\left[\widehat{S} \leq k\right] = \mathrm{F}_{\widehat{S}}(k) = \begin{cases} \dfrac{\alpha-1}{\mathsf{m}\alpha} \cdot k & \forall k \leq \mathsf{m} \\ 1 - \dfrac{1}{\alpha} \cdot \left(\dfrac{\mathsf{m}}{k}\right)^{\alpha-1} & \forall k \geq \mathsf{m}. \end{cases}$$

APPENDIX B: IMPLEMENTATION NOTES

We have developed the C++ class `Cox`, following the interface of the class `Random` of the *GNU libg++* library as a guideline. We have attempted to provide an implementation as efficient as possible, intending to have approximately the same level of efficiency as in a generator of any random variable. The source code is available at `<ftp://ftp-gris.det.uvigo.es/pub/LRD/MGinf-src.tgz>`.

An object of class `Cox` has the following member objects:

- `servers` An object of class `ListTimes` which stores the number of users in the system, their departure times (in slots) and the isochronous clock.
- `batches` An array of K objects of class `IntPoisson` which generate samples of Poisson random variables with mean values $\lambda \cdot \Pr[S=i]$ for $i = 1, 2, \ldots, \mathsf{K}$.
- `batch` An object of class `IntPoisson` which generates samples of a Poisson random variable with mean value $\lambda \cdot \Pr\left[S > \mathsf{K}\right]$.
- `demand` An object of class `IntPareto_U_cond` which generates samples of the random variable $S|_{S>\mathsf{K}}$.

IntRandom

Both classes `IntPareto_U_cond` and `IntPoisson` are built from the class `IntRandom`, which implements a generic tabular method to invert the cdf $\mathrm{F}(k)$ of a non-negative discrete random variable (Suárez et al. 2002). In this way, the efficiency of the `IntPoisson` generator is almost independent of the mean value λ.

ListTimes

This class stores the number of users in the M/G/∞ system along with their departure times. Since the simulation clock is isochronous, we need a set of counters, each one of them storing the number of departures in a given future time $t = n$.

The straightforward way to manage these counters is through a single-linked list. To process an arrival, we have to locate in the list the node corresponding to the new user's departure time, and increase its counter; or, if it does not yet exist, to insert a new node in the list in the correct place.

Given that the search in the list for every arrival is costly, we complete this data structure with a cyclic array that stores those counters with departure time within the next V units of time. Thus, if the service time of a new user is less than V, the counter associated to its departure time is directly accessed. By selecting V as a power of two the code to insert an arrival is:

```
inline void ListTimes::In
(unsigned long service_time,unsigned num)
{
  inside+=num;
  if (service_time < V)
    next_outs[(idx+service_time) & Vmask]+=
      num;
    // Vmask is V-1, and idx is clock & Vmask;
  else
    InList(service_time,num);
}
```

where `InList` inserts `num` arrivals in the single-linked list `out_jobs` in the usual way.

The method to advance and generate the sample is implemented as:

```
inline unsigned ListTimes::Tick()
{
  inside -= next_outs[idx];
  if (out_jobs && out_jobs->o_time ==
    clock+V)
  {
    Out_jobs *aux = out_jobs;
    out_jobs = out_jobs->next;
    next_outs[idx] = aux->number;
    delete aux;
  }
  else
    next_outs[idx] = 0;
  clock++;
  idx = clock & Vmask;
  return inside;
}
```

REFERENCES

Cox, D. R. 1984. Long range dependence: A review. *Statistics: An Appraisal*. Iowa State University Press.

Cox, D. R and V. Isham. 1980. *Point Processes*. Chapman and Hall.

Crouse, M. S. and R. G. Baraniuk. 1999. *Fast, exact synthesis of Gaussian and nonGaussian long-range-dependent processes*. Draft version available at `<http://www.dsp.rice.edu/publications/pub/fftMC.ps.gz>`.

Erramilli, A., O. Narayan and W. Willinger. 1996. Experimental queueing analysis with long-range dependent packet traffic. *IEEE/ACM Transactions on Networking* 4: 209–223.

Garrett, M. W. and W. Willinger. 1994. Analysis, modeling and generation of self-similar VBR video traffic. In *Proceedings of ACM SIGCOMM'94*. London, UK.

Grossglauser, M. and J. C. Bolot. 1996. On the relevance of long-range dependence in network traffic. In *Proceedings of ACM SIGCOMM'96*.

Hurst, H. E. 1951. Long-term storage capacity of reservoirs. *Trans. Am. Soc. Civil Engineers* 116: 770–799.

Krunz, M. M. and A. M. Makowski. Modeling Video Traffic Using M/G/∞ Input Processes: A Compromise Between Markovian and LRD Models. *IEEE Journal on Selected Areas in Communications* 16 (5): 733–748.

Leland, W. E., M. S. Taqqu, W. Willinger, and D.V. Wilson. 1994. On the self-similar nature of Ethernet traffic (extended version). *IEEE/ACM Transactions on Networking* 2 (1): 1–15.

Likhanov, N., B. Tsybakov, and N. Georganas. 1995. Analysis of an ATM buffer with self-similar ("fractal") input traffic. In *Proceedings of IEEE INFOCOMM'95*. Boston, MA.

Parulekar, M. and A. M. Makowski. 1996. Tail probabilities for a multiplexer with self-similar traffic. In *Proceedings of IEEE INFOCOMM'96*. San Francisco, CA.

Poon, W. and K. Lo 2001. A refined version of M/G/∞ processes for modeling VBR video traffic. *Computer Communications* 24 (11): 1105–1114.

Suárez, A., J. C. López, C. López, M. Fernández, R. Rodríguez, and M. E. Sousa. 2002. A new heavy-tailed discrete distribution for LRD M/G/∞ sample generation. *Performance Evaluation* 47 (2/3): 197–219.

Tsybakov, B. and N. D. Georganas. 1997. On self-similar traffic in ATM queues: Definitions, overflow probability bound, and cell delay distribution. *IEEE/ACM Transactions on Networking* 5 (3): 397–409.

AUTHOR BIOGRAPHIES

MARÍA ESTRELLA SOUSA-VIEIRA is an assistant professor in the *Departamento de Enxeñería Telemática* at *Universidade de Vigo*. She received a telecommunication engineering degree from *Universidade de Vigo* in 1998. Since then, she is working towards her Ph.D. in the *Departamento de Enxeñería Telemática* at *Universidade de Vigo*. Her e-mail address is `<msousa@det.uvigo.es>`.

ANDRÉS SUÁREZ-GONZÁLEZ is an associate professor in the *Departamento de Enxeñería Telemática* at *Universidade de Vigo*. He received a Ph.D. degree in telecommunication engineering from *Universidade de Vigo* in 2000. He is a member of ACM. His current research interests include simulation methodology and analysis of stochastic systems.

CÁNDIDO LÓPEZ-GARCÍA is an associate professor in the *Departamento de Enxeñería Telemática* at *Universidade de Vigo*. He received a Ph.D. degree in telecommunication engineering from *Universidad Politécnica de Madrid* in 1995. His main research interests are in the area of performance evaluation in communication networks.

MANUEL FERNÁNDEZ-VEIGA is an associate professor in the *Departamento de Enxeñería Telemática* at *Universidade de Vigo*. He received a Ph.D. degree in telecommunication engineering from *Universidade de Vigo* in 2001. He is an associate member of IEEE.

JOSÉ C. LÓPEZ-ARDAO is an associate professor in the *Departamento de Enxeñería Telemática* at *Universidade de Vigo*. He received a Ph.D. degree in telecommunication engineering from *Universidade de Vigo* in 1999. He is a member of ACM.