ELSEVIER

Advances in Engineering Software 37 (2006) 75–84

ADVANCES IN ENGINEERING SOFTWARE

www.elsevier.com/locate/advengsoft

Generating correlated matrix exponential random variables☆

S. Fitzgerald[a], J. Place[b,*], A. van de Liefvoort[b]

[a]*Department of Information and Computer Sciences, Metropolitan State University, 730 Hennepin Ave., Minneapolis, MN 55403-1897, USA*
[b]*School of Computing and Engineering, University of Missouri—Kansas City, 5100 Rockhill Rd., Kansas City, MO 64110, USA*

Received 12 June 2004; received in revised form 18 April 2005; accepted 21 April 2005
Available online 15 June 2005

Abstract

In this paper, we focus on inter-arrival time autocorrelation and its impact on model performance. We present a technique to generate matrix exponential random variables that match first-order statistics (moments) and second-order statistics (autocorrelation) from an empirical distribution. We briefly explain the matrix exponential distribution and show that we can represent any empirical distribution arbitrarily closely as matrix exponential. We then show how we can incorporate an autocorrelation structure into our matrix exponential random variables using the *autoregressive to anything* technique. We present examples showing how we match first and second-order statistics from empirical distributions and finally we show that our autocorrelation matrix exponential random variables produce more accurate performance metrics from simulation models than traditional techniques.

Keywords: Random variables; Teletraffic models; Autocorrelation linear algebra queueing theory; Simulation

1. Introduction

There is an ever growing literature on the generation of IID random numbers. See for example, the latest edition of Knuth, vol. II [1] or on-line see http://random.mat.sbg.ac.at/links/. For correlated random variables, the literature is much smaller, with most of the results being for pairwise correlation, or autocorrelation for a specific family of distributions. There are only a few attempts to generate autocorrelated random numbers with any marginal distribution. Polge's method [2] generated first a list of i.i.d. random variates, which is then sorted in such a way as to induce the desired correlation. This makes the method expensive in both time and space if large numbers of RV's are to be generated. In 1979, Badel [3] developed a method for the extended exponential family, based on a first-order autoregressive scheme with an additive Bernoulli residual. Metzner [4] generated autocorrelated numbers where the desired marginal is approximated with a four-moment matched distribution (even for the 'simple' distributions) and the allowed correlations are restricted. More recently, Willemain and Desautels [5] presented the 'Sum of Uniforms' method to generate correlated uniform random variables. This is a simple and elegant method, but limited to serial correlation $AR(1)$ only. Myers and Yeh [6] present a simple case of generating discrete RV's from a specified autocorrelation structure. Parkinson [7] discusses generating RV's from curves derived from second-order linear segments of an autocorrelation structure. Melamed et al. [8–12] developed the TES-family of correlated random numbers in a series of papers, but the process is rather counter intuitive and the software package TEStool is reportedly no longer available. More recently, Cario and Nelson [13–15] presented the ARTA method (Auto-Regressive-To-Anything). This method uses an underlying $AR(p)$ Gaussian autoregressive base process which is transformed into the desired marginals. See also [16, pp. 66–71]. It is of critical importance here to choose correctly the autocorrelation of the base process, such that the transformed process has the desired autocorrelation. We use this method, which we pair with the generation of ME random variables.

Generating RV's from known distributions, e.g. Poisson, exponential, Weibull, is well understood and well documented [17, Chapt. 8]. Suppose, however, that RV's are

☆ This work was partially supported under NSF Grant CDA-9422092.
* Corresponding author.
E-mail addresses: sne.fitzgerald@metrostate.edu (S. Fitzgerald), place@umkc.edu (J. Place).

0965-9978/$ - see front matter
doi:10.1016/j.advengsoft.2005.04.006

desired from an empirical distribution, i.e. a sample whose cumulative distribution function (CDF) $F(x)$ is unknown. The simulation engineer has three broad choices. Find the parameters within a familiar distribution such that the CDF of the familiar distribution 'fits' the target distribution and then generate RV's from this target distribution [17, p. 329] or one may generate RV's from a frequency table representing the CDF [17, p. 494] of the empirical distribution or bootstrapping techniques can be used to generate RV's [18–20].

Fitting empirical data to a family of known distributions such as the exponential often misrepresents important parts of the empirical distribution—particularly the head and tail regions. Generating RV's from a table representing the empirical distribution function often results in modeling error due to the inability of the table to accurately reflect the shape of the distribution in the tail or other low-probability region and bootstrapping also may be unsuitable for estimating the empirical distribution in the tail distribution [21].

Brown et al. [22,23] demonstrated that RV's can be generated when the CDF $F(x)$ is known both analytically and empirically. Brown's techniques use a matrix exponential representation of the target distribution and then generate matrix exponential random variables and produce RV's that are based on first-order statistics (moments) of the target distribution. There are cases, however, where RV's based on first-order statistics are not powerful enough. Teletraffic models are a good example.

Buffering requirements at a network node or switch in an ATM network is an important aspect of telecommunication network design. It has been shown that models that assume a renewal (uncorrelated) arrival stream, i.e. uncorrelated inter-arrival times, result in operational network buffer designs which frequently exhibit much lower cell loss than actual traffic conditions [24–26]. Analysis of network traffic shows that packet/cell inter-arrival times are often highly autocorrelated.

In 1989 and 1990, Leland and Wilson [27] collected traces of several hundred million Ethernet packets from Bellcore's Morristown Research and Training Center. These packet traces formed the basis of the conclusion in several papers [28–32] that renewal models of network traffic may poorly reflect the behavior of communication networks. The behavior observed by Leland et al. was earlier characterized by Mandelbrot [33] as *self-similar*.

In their landmark work, Leland et al. [28,29] show that the Ethernet network traffic was self-similar across widely varying time scales. Traffic bursts had no natural length and at every time scale, periods of activity are clearly distinguishable from idle times. The principal characteristic of self-similar distributions and the basis for the characterization of these distributions as self-similar is long-range dependency. Cell inter-arrival times have been found to be autocorrelated at lags of 5000 cells. This long-term dependency is why traditional Poisson and other renewal streams are not viable models of these systems. Short and long-term dependencies have also been observed in TCP applications over wide-area networks [34,25] and VBR video traffic has been shown to have strong correlations [35]. These correlations cause buffer overow which is not predicted in models based on renewal input streams.

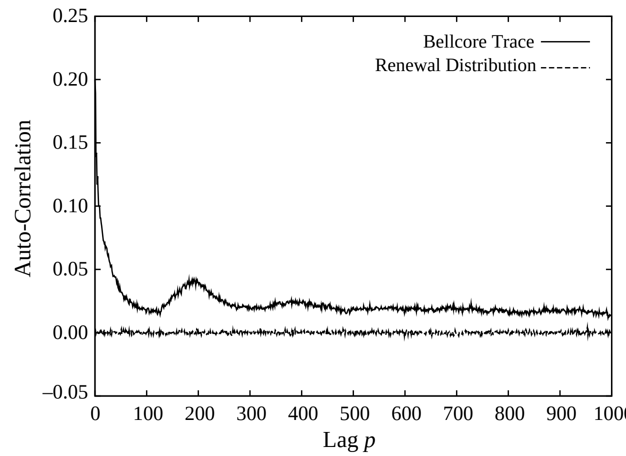

Fig. 1. Inter-arrival time autocorrelation comparison.

Fig. 1 shows the packet inter-arrival time autocorrelation in one of the Bellcore Ethernet traces.[1] Packet inter-arrival time autocorrelation is plotted to lag-1000. The inter-arrival time autocorrelation for an exponential random variable with the same mean as the Bellcore trace is also plotted to lag-1000 to allow comparison.

Note the slow decay in packet inter-arrival time autocorrelation for the Bellcore trace as the lag increases while the exponential inter-arrival times have essentially no autocorrelation at *any* lag-p. Clearly, a simulation model with a Poisson arrival stream will not be an effective surrogate for a system with autocorrelated arrivals such as those found in the Bellcore Ethernet packet trace even if the Poisson stream matches first-order statistics of the Bellcore trace.

In this paper, we present a technique that will allow the simulation engineer to generate correlated RV's, that not only match first-order statistics from the target CDF arbitrarily closely, but also match second-order statistics from the target CDF arbitrarily closely. This technique captures both the first and second-order behavior through mathematical formulations which are then used to generate the RV stream. Thus, the simulation engineer can generate autocorrelated RV's that match *any* distribution arbitrarily closely.

Our technique creates a matrix exponential (ME) representation of the target CDF that matches first-order statistics arbitrarily closely then we overlay our ME RV's with *any* correlation structure. As Brown [22,23] has shown, this is not a 'distribution fitting' technique because we

[1] August 1990 Ethernet trace.