MODELING AND SIMULATING TIME SERIES INPUT PROCESSES WITH ARTAFACTS AND ARTAGEN

Marne C. Cario

Delphi Packard Electric Systems
P.O. Box 431
Warren, Ohio 44486, U.S.A.

ABSTRACT

We develop an efficient numerical method for fitting ARTA processes for use as simulation input. ARTA processes are stationary time series with arbitrary marginal distributions and autocorrelations specified through finite lag p . We discuss the software package **ARTAFACTS**, which implements the numerical method, and the package **ARTAGEN**, which generates observations from ARTA processes. To demonstrate the use of the software, we present a real-world example.

1 INTRODUCTION

Many systems contain processes that are inherently dependent. For instance, the temperatures in an oven of a steel mill, measured over time, are likely to be autocorrelated. The sizes of demands on an inventory system are often dependent; for instance, a small demand may follow a very large demand. Daily reserve levels in a blood bank are autocorrelated. The prevalence of dependent input processes, as well as the potential error in the simulation output if dependencies are ignored (see, for example, Livny, Melamed, and Tsiolis 1993), motivates the need for simulation input models that represent autocorrelated time series.

ARTA (AutoRegressive To Anything), introduced in Cario and Nelson (1996), is a transformation-oriented approach for modeling and generating a time-series process $\{Y_t; t = 1, 2, \dots\}$ with an arbitrary marginal distribution and autocorrelation structure specified through lag p . This approach takes a process with a known autocorrelation structure, the base process $\{Z_t\}$, and transforms it to achieve the desired marginal distribution for the input process, $\{Y_t\}$. The target autocorrelation structure of $\{Y_t\}$ is obtained by adjusting the autocorrelation structure of the base process. In our model, the base process is a standardized Gaussian autoregressive process of order p ($AR(p)$). Thus, the major step in constructing

an ARTA process is finding the autocorrelation structure of the $AR(p)$ process that results in the desired autocorrelation structure for the ARTA process. We present an efficient numerical procedure for performing this step.

The paper is organized as follows: Section 2 discusses other methods for generating time-series inputs for simulation. Section 3 gives a brief review of ARTA processes. In Section 4 we present a specialized numerical method for finding the required autocorrelation structure of the $AR(p)$ process. We describe **ARTAFACTS** (ARTA Fitting Algorithm for Constructing Time Series), a Fortran implementation of the numerical procedure, in Section 5. In Section 6 we discuss the software package **ARTAGEN** (AutoRegressive To Anything GENERation), which generates observations from an ARTA process for use as simulation input. Section 7 provides a real-world application of an ARTA process. Conclusions appear in Section 8.

2 BACKGROUND

A variety of methods exist for modeling time-series input processes with specified marginal distributions and autocorrelation structures. These methods differ in their degree of generality and their ease of use.

In Song, Hsiao, and Chen (1996), the authors use a transformation-oriented approach with an $AR(1)$ base process. However, they attempt to match only the lag-1 autocorrelation, and they use simulation to find the lag-1 autocorrelation of the base process that gives the desired lag-1 autocorrelation of the input process. This approach becomes computationally prohibitive if extended to more than two or three autocorrelations. The ARTA approach matches $p \geq 1$ autocorrelations, and is computationally more efficient.

TES (Melamed 1991) also uses a transformation-oriented approach to modeling time-series input pro-

cesses. TES, which is implemented in the software package **TESTool** (Melamed, Hill, and Goldsman 1992), uses a series of autocorrelated uniform random variates as the base process. To obtain an input process with the desired characteristics using **TESTool**, the user must adjust the distribution of the base-process noise term. Unfortunately, there is no standard procedure for adjusting the noise term, so constructing a TES process is an art as well as a science. Furthermore, TES processes often exhibit unusual sample paths. The sample paths can be improved at the expense of further complicating the modeling process.

A completely different approach, for which there is a large literature, is to construct a time-series process that exploits properties that are specific to the particular marginal distribution of interest for $\{Y_t\}$. An example is Lewis, McKenzie and Hugus (1989) who construct time series with gamma marginals. The primary shortcoming of this type of approach is that it is not general: a different model is required for each marginal distribution of interest. In addition, the sample paths of these processes, while adhering to the desired marginal distribution and autocorrelation structure, sometimes have unexpected features.

3 ARTA PROCESSES

This section reviews the definition of an ARTA process and summarizes the results of Cario and Nelson (1996). Formally, an ARTA process is a model of a stationary time series $\{Y_t\}$ with the following properties:

1. $Y_t \sim F_Y, t = 1, 2, \dots$, where F_Y is an arbitrary cumulative distribution function (cdf); and
2. $(\rho_1, \rho_2, \dots, \rho_p) = \boldsymbol{\rho} = (\text{Corr}[Y_t, Y_{t+1}], \text{Corr}[Y_t, Y_{t+2}], \dots, \text{Corr}[Y_t, Y_{t+p}])$

where F_Y and $\boldsymbol{\rho}$ are given. In an ARTA process, $\{Y_t\}$ is represented as a transformation of a standardized Gaussian AR(p) process, as follows:

ARTA Process

1. Let $\{Z_t; t = 1, 2, \dots\}$ be a stationary Gaussian AR(p) process

$$Z_t = \alpha_1 Z_{t-1} + \alpha_2 Z_{t-2} + \dots + \alpha_p Z_{t-p} + \varepsilon_t,$$

where $\{\varepsilon_t\}$ is a series of independent $N(0, \sigma^2)$ random variables and σ^2 is selected so that the marginal distribution of the $\{Z_t\}$ process is $N(0, 1)$, where N denotes the normal distribution.

2. Define the ARTA process $Y_t = F_Y^{-1}[\Phi(Z_t)]$, $t = 1, 2, \dots$, where Φ is the standard normal cdf.

The transformation $F_Y^{-1}[\Phi(\cdot)]$ ensures that $\{Y_t\}$ has the desired marginal distribution F_Y . Therefore, the central problem is to select the autocorrelation structure $\mathbf{r} = (r_1, r_2, \dots, r_p)$ for the AR(p) process, $\{Z_t\}$, that gives the desired autocorrelation structure $\boldsymbol{\rho}$ for the input process, $\{Y_t\}$.

The autocorrelation structure of the AR(p) base process directly determines the autocorrelation structure of the input process, since $\text{Corr}[Y_t, Y_{t+h}] = \text{Corr}\{F_Y^{-1}[\Phi(Z_t)], F_Y^{-1}[\Phi(Z_{t+h})]\}$. To adjust this correlation, we can restrict attention to adjusting $E[Y_t Y_{t+h}]$, since

$$\text{Corr}[Y_t, Y_{t+h}] = \frac{E[Y_t Y_{t+h}] - (E[Y_t])^2}{\text{Var}[Y_t]},$$

and $E[Y_t]$ and $\text{Var}[Y_t]$ are fixed by F_Y . Then, since (Z_t, Z_{t+h}) has a standard bivariate normal distribution with $\text{Corr}[Z_t, Z_{t+h}] = r_h$, we have

$$\begin{aligned} E[Y_t Y_{t+h}] &= E\{F_Y^{-1}[\Phi(Z_t)] F_Y^{-1}[\Phi(Z_{t+h})]\} \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} F_Y^{-1}[\Phi(z_t)] F_Y^{-1}[\Phi(z_{t+h})] \\ &\quad \times \varphi_{r_h}(z_t, z_{t+h}) dz_t dz_{t+h}, \end{aligned} \quad (1)$$

where φ_{r_h} is the standard bivariate normal probability density function (pdf) with correlation r_h . We are only interested in processes for which this expectation exists.

Observe from Equation (1) that the lag- h autocorrelation of $\{Y_t\}$ is a function only of the lag- h autocorrelation of $\{Z_t\}$, which appears in the expression for φ_{r_h} . We denote this function by $\rho(r_h)$. Thus, the problem of determining the autocorrelations for $\{Z_t\}$ that give the desired autocorrelations for $\{Y_t\}$ reduces to p independent problems: For each lag $h = 1, 2, \dots, p$, find the value r_h for which $\rho(r_h) = \rho_h$.

Cario and Nelson (1996) prove that for each lag h all feasible values of ρ_h are attainable under ARTA processes. By feasible, we mean that $\underline{\rho} \leq \rho_h \leq 1$, where

$$\underline{\rho} = \frac{\int_0^1 F_Y^{-1}(u) F_Y^{-1}(1-u) du - (E[Y_t])^2}{\text{Var}[Y_t]}$$

is the minimum feasible bivariate correlation for F_Y (see Whitt 1976).

In general, it is not possible to find the r_h -values analytically; therefore, we develop an efficient numerical procedure to find the r_h -values to within any precision. Cario and Nelson (1996) prove that $\rho(r)$ is nondecreasing and, under very mild conditions on F_Y , it is also continuous. These results guarantee that our numerical procedure will converge.

4 NUMERICAL PROCEDURE

There are two main considerations in constructing an efficient numerical procedure for finding the autocorrelation structure of $\{Z_t\}$ which gives the target autocorrelation structure for $\{Y_t\}$. The first consideration is the choice of a method to integrate Equation (1) numerically. The second consideration is the development of a numerical search procedure that selects and updates the values of r_h in an efficient manner.

Several properties of the integrand make it non-trivial to integrate Equation (1) numerically. Foremost is the fact that the integrand is bivariate, and there are few general-purpose numerical-integration methods for multivariate integrals. Our application requires a general-purpose method since we do not make any restrictions on the functional form of the marginal distribution F_Y other than those required for it to be a cdf and for the integral (1) to exist. For instance, we do not require F_Y to be continuous. The doubly infinite ranges of the inner and outer integrals further complicate the integration since many numerical-integration methods for multivariate integrals apply to subspaces with special shapes, such as spheres and cubes. Finally, since Equation (1) must be calculated for many values of r_h , and it is often computationally expensive to calculate F_Y^{-1} , we need a method that efficiently manages the number of function evaluations.

Taking the previously mentioned factors into consideration, our solution is to iteratively apply a univariate numerical-integration method to compute Equation (1). Specifically, we use a bivariate implementation of the Gauss-Kronrod quadrature rule (Piessens et al. 1983, pp. 16–17). Dropping the time subscripts t and $t + h$ from z , the Gauss-Kronrod quadrature rule approximates Equation (1) as:

$$\begin{aligned} \widehat{E}[Y_t Y_{t+h}] &= \sum_{i=1}^{15} \sum_{j=1}^{15} w(z_i) w(z_j) F_Y^{-1}[\Phi(z_i)] \\ &\quad \times F_Y^{-1}[\Phi(z_j)] \varphi_{r_h}(z_i, z_j), \end{aligned} \quad (2)$$

where $w(z)$ is a weight function and the z_i s and z_j s are abscissae. Piessens et al. (1983) give the abscissae for the interval $(-1, 1)$, but they may easily be mapped onto the range of interest. In our implementation, we truncate each doubly infinite range to the finite range $(-10, 10)$, and use a linear mapping of the abscissae from $(-1, 1)$ to $(-10, 10)$. Using the finite range does not decrease the accuracy of the approximation significantly since the value of the bivariate normal pdf is effectively zero outside of this range.

We set absolute-error tolerances for the inner and outer integral using the method given on pp. 112–

113 of Piessens et al. (1983). The absolute error is taken to be the absolute error of the difference between two sums: the Kronrod sum, for which all of the weights are non-zero, and the Gauss sum, for which some of the weights are zero. If either the inner or the outer integral exceeds its absolute-error tolerance on the initial integration range $(-10, 10)$, then the corresponding integration range is bisected and the integral is computed on each half-interval. This procedure is repeated until all of the inner and outer integrals meet the absolute-error tolerances, giving us a bound on the overall integration error as well.

Given a value of r , we may compute the value of $\rho(r)$ using the Gauss-Kronrod quadrature rule. The remaining step is to construct an efficient numerical search procedure to find the autocorrelation structure $\mathbf{r}^* = (r_1^*, r_2^*, \dots, r_p^*)$ of $\{Z_t\}$, such that $\boldsymbol{\rho}^* = (\rho(r_1^*), \rho(r_2^*), \dots, \rho(r_p^*)) \approx \boldsymbol{\rho}$. Many distributions have special properties which enhance the efficiency of such a numerical search procedure.

When F_Y is a continuous uniform distribution, a closed-form solution for $\rho(r)$ is given in Li and Hammond (1975, Equation 7) as:

$$\rho(r) = 2 \sin \frac{\pi}{6} r.$$

Furthermore, for some distributions—such as the exponential distribution with rate parameter λ , or any other scale family— $\rho(r)$ is independent of the distribution parameters. To increase the efficiency of the search procedure when $\rho(r)$ is parameter-independent, we access a precomputed table of $(r, \rho(r))$ values to find a good starting point for the search. For symmetric distributions, $\rho(-r) = -\rho(r)$ (Song, Hsiao, and Chen 1996); thus, it is only necessary to compute Equation (1) for positive values of r . Our numerical search procedure makes use of these results.

When Y_t is a discrete random variable with finite support, we use a special-purpose numerical-integration method. In this case, which includes the empirical cdf, it is easy to show that Equation (1) is a finite sum of terms of the form

(constant) $\times$ (area under the bivariate normal pdf).

In general, the form of Equation (2) leads to additional computational savings. Observe that in Equation (2), φ_{r_h} is the only term that is a function of r_h . Since for many distributions calculating F_Y^{-1} is computationally intensive, we compute Equation (2) for many values of r_h simultaneously. Specifically, at each pair of abscissae (z_i, z_j) , we calculate $w(z_i)w(z_j)F_Y^{-1}[\Phi(z_i)]F_Y^{-1}[\Phi(z_j)]$ once, followed by $\varphi_{r_h}(z_i, z_j)$ for several values of r_h . Then, after

multiplying the two terms together, we keep track of the Gauss and the Kronrod sums for each value of r_h separately. The error for each integration interval is taken to be the maximum absolute error over all values of r_h , and the interval is bisected if the error does not meet the given tolerance.

In brief, our numerical search procedure selects a grid of r_h -values at each iteration. At least initially, these grid points depend upon F_Y and ρ . Then, the Gauss-Kronrod quadrature rule is used to integrate Equation (1) simultaneously for all of the grid points. For each autocorrelation lag h , the procedure finds the two grid points r_{j_h} and r_{j_h+1} such that $\rho(r_{j_h}) \leq \rho_h \leq \rho(r_{j_h+1})$. If one of these two grid points meets the relative-error criterion, then the search is complete for lag h . Otherwise, several points between r_{j_h} and r_{j_h+1} are added to the grid for the next iteration. The procedure is repeated until the relative-error criteria are met for all p lags.

5 ARTAFACFS—SOFTWARE FOR FITTING ARTA PROCESSES

ARTAFACFS is a Fortran implementation of the numerical search procedure outlined in Section 4. User input, which is read from a file, includes the marginal distribution F_Y , the number of autocorrelation lags p to match, the autocorrelations ρ to match, the relative-error criterion for each autocorrelation lag, the parameters of the ARTA-process marginal distribution, and the name of the output file. The following marginal distributions are supported: Normal, Student's t , Continuous Uniform, Exponential, Gamma, Weibull, Lognormal, Johnson Unbounded, Discrete with finite support and Empirical. The output consists of the first p autocorrelation lags of the $AR(p)$ process, along with the resultant autocorrelation values for the ARTA process, which are within the user-specified relative errors. The output also includes the AR parameters $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_p)$.

If necessary, ARTAFACFS can read in time-series data from a file, calculate the sample autocorrelation function of the series and test whether the autocorrelations are significantly different from 0. For this option, the user specifies the number of autocorrelation lags to calculate and the name of the file to which the sample autocorrelation structure will be written.

Prior to invoking the numerical search procedure, ARTAFACFS performs several checks on the feasibility of the user input. For instance, ARTAFACFS checks that the correlation matrix implied by the autocorrelation structure for the input process is positive definite. Furthermore, ARTAFACFS checks that $\underline{\rho} \leq \rho_h \leq 1, h = 1, 2, \dots, p$, where $\underline{\rho}$, the minimum

feasible bivariate correlation for F_Y , is calculated by ARTAFACFS. These two conditions are necessary, but not sufficient, for the ARTA process to be stationary. After the numerical procedure terminates, ARTAFACFS determines whether or not the underlying $AR(p)$ process is stationary. If the $AR(p)$ process is not stationary, a warning appears in the output file and no ARTA representation is possible.

Figure 1 displays a sample ARTAFACFS input file for an application that will be presented in the next section. The first line contains the name of the output file, `weib.out`. The 0 on the second line indicates that we do not want ARTAFACFS to read in an empirical time series and estimate its autocorrelation structure. Instead, the third line tells ARTAFACFS that we want the marginal distribution of the ARTA process to be a Weibull distribution, which has code number 6. The 3 on the next line indicates that we want to match three autocorrelation lags of the input process. The following six lines give the desired autocorrelations and their corresponding relative-error criteria for lags 1 through 3. The final two lines of input contain the scale parameter, 0.9432, and shape parameter, 5.14, of the Weibull distribution.

```
weib.out
0
6
3
0.749
0.01
0.409
0.01
0.121
0.01
0.9432
5.14
```

Figure 1: Sample ARTAFACFS Input File, `input.dat`

Figure 2 displays the ARTAFACFS output file `weib.out` corresponding to the input file given in Figure 1. The output file reiterates that the ARTA distribution is a Weibull(0.9432, 5.14) distribution and that we want to match three autocorrelation lags. The output file also gives the desired autocorrelations and their required relative errors. Next, it displays three columns of program output, the three ARTA-process autocorrelations $\rho^* = (0.751, 0.407, 0.121)$, their underlying AR-process autocorrelations $\mathbf{r}^* = (0.752, 0.408, 0.122)$ such that $\rho(r_h^*) \approx \rho_h$, for $h = 1, 2, 3$, and the actual relative errors that were attained. The final column of output contains the pa-

parameters α_1, α_2 , and α_3 of the AR(3) base process.

```
***** USER INPUT *****

ARTA Distribution: Weibull
Scale Parameter:   0.94320
Shape Parameter:   5.14000

Number of ARTA autocorrelation lags to match:  3

Lag   Desired Autocorrelation   Desired Relative Error
***   *****
  1     0.7490000000000000      0.01000
  2     0.4090000000000000      0.01000
  3     0.1210000000000000      0.01000

Minimum feasible bivariate correlation is: -0.98875

*****

***** ARTA PROGRAM OUTPUT *****

Lag      ARTA Correl.      AR Correl.      Rel. Error

  1      0.750955          0.752000          0.002610
  2      0.406725          0.408000          -0.005563
  3      0.121049          0.121600          0.000402

i          Alpha(i)
1           1.0017
2          -0.2978
3          -0.0632
```

Figure 2: Sample ARTAFACTS Output File

6 ARTAGEN—SOFTWARE FOR GENERATING ARTA PROCESSES

After fitting an ARTA process using ARTAFACTS, we can use the AR output parameters α to generate simulation inputs (see Cario and Nelson 1996). ARTAFACTS automatically writes the input file for ARTAGEN, a software package that generates observations from ARTA processes.

ARTAGEN reads input from a file called `input.gen`. A sample input file for ARTAGEN appears in Figure 3. The first line of the input file specifies the name of the file, `weib.gen`, to which the generated observations will be written. The second line contains the name of the file, `weib.sum`, to which summary statistics about the generated observations will be written. The next three lines contain the code number (6) of the ARTA marginal distribution, the number of autocorrelation lags that have been matched (3), and the number of observations to generate (1000). The au-

tocorrelations for the AR(p) process followed by the autocorrelations for the ARTA process appear on the next six lines. The final two lines contain the parameters of the marginal distribution.

```
weib.gen
weib.sum
6
3
1000
0.752
0.408
0.1216
0.749
0.409
0.121
0.9432
5.14
```

Figure 3: Sample ARTAGEN Input File, `input.gen`

7 APPLICATION

Continuous-flow production lines, such as those used to extrude plastics, are common in the chemical industry. Process variables such as temperatures and pressures are often key parameters on these types of production lines, and understanding their effects on the manufacturing system is critical. Systems simulation can be used to model new and existing continuous-flow lines, as well as to train new operators in proper responses to process changes.

In this section, we present time-series data for a pressure variable from a continuous-flow polymer process at a manufacturing plant of a large chemical company. We fit a marginal distribution using the commercially available ARENA input processor and estimated the autocorrelation structure using ARTAFACTS. We then used ARTAFACTS to fit an ARTA process to the data set. Finally, we generated 1000 observations from the ARTA process using ARTAGEN. To assess the fit of the ARTA process, we compared the properties of the empirical time series with those of the generated time series.

Figure 4 displays the time-series plot of the empirical data. The sample contains 519 observations, with sample average $\bar{\mu} = 45.48$ and sample variance $s^2 = 0.039$. The sample autocorrelation function through lag 3 is $\hat{\rho} = (0.749, 0.409, 0.121)$. No other autocorrelations are statistically different from zero.

ARENA selected $44.6 + \text{Weibull}(0.9432, 5.14)$ as the best fit for the marginal distribution of the pressure

Figure 4: Time-series Plot of the Empirical Data

variable. We used this Weibull distribution (minus the constant) and $\hat{\rho}$ as the input to ARTAFACTS, which is displayed in Figure 1. Recall from the output file in Figure 2 that the autocorrelation structure for the underlying AR(3) process was $\mathbf{r}^* = (0.752, 0.408, 0.122)$ and the implied ARTA-process autocorrelation structure was $\hat{\rho}^* = (0.751, 0.407, 0.121)$.

Given $\mathbf{r}^*$, we used ARTAGEN to generate an ARTA process with the specified Weibull distribution and autocorrelation structure $\hat{\rho}^* \approx \rho$. Recall that the ARTAGEN input file for this process appears in Figure 3. Figure 5 contains the ARTAGEN summary statistics file.

As indicated in Figure 5, the sample average for this series is $\bar{\mu} = 44.6 + 0.86 = 45.46$ and the sample variance is $s^2 = 0.0387$. The sample autocorrelation function through lag 3 is $\hat{\rho} = (0.739, 0.383, 0.081)$. The differences between the sample statistics and the theoretical values are due to sampling error—they would converge to the theoretical values as the number of generated observations increased.

Figure 6 displays the time-series plot of the observations from the ARTA process that we fit to the pressure variable. Comparing Figures 4 and 6, observe that the sample paths are qualitatively similar, although the height of the spikes is more variable in the empirical time series, and the ARTA process varies more consistently about its mean than does the empirical time series. The differences between the two time series reflect sampling error as well as the fact that the marginal distribution and autocorrelation structure do not capture all of the characteristics of a time-series process.

The scatterplots of (Y_t, Y_{t+1}) for the empirical and ARTA data on the pressure variable appear in Figures 7 and 8, respectively. Observe that the ARTA data

***** USER INPUT *****

ARTA Distribution: Weibull
Scale Parameter: 0.94320
Shape Parameter: 5.14000

Lag	ARTA Corr.	AR Corr.
1	0.7490	0.7520
2	0.4090	0.4080
3	0.1210	0.1216

**** SUMMARY STATISTICS ****

Average of Output (Y) Process: 0.86158
Sample Variance of Output (Y) Process: 0.03869

Lag	Correlation of Output (Y) Process
1	0.73937780
2	0.38340840
3	0.08111539

Average of Input (Z) Process: -0.02773
Sample Variance of Input (Z) Process: 1.02594

Lag	Correlation of Input (Z) Process
1	0.74147892
2	0.38993248
3	0.08886040

Figure 5: Sample ARTAGEN Summary File, weib.sum

is more scattered than the empirical data, and that there are a few more points in the bottom, left corner of the scatterplot for the empirical data. Similar differences appear in the scatterplots for lags two and three, which are not displayed. The differences are compatible with the differences between the fitted distribution and the empirical distribution, along with the differences between the time-series plots. Overall, the ARTA process provides a plausible model for the empirical time series.

8 CONCLUSIONS

The example in Section 7 demonstrates the usefulness of ARTA processes for modeling time-series inputs for simulation. Certainly, they provide a more faithful representation of the empirical systems than do series of independent, identically distributed random variables, which are often used in practice. Furthermore, the software packages ARTAFACTS and ARTAGEN simplify the tasks of building and generating observations from ARTA models. Instructions for downloading and using ARTAFACTS and ARTAGEN appear on the ARTA web page at <http://primal.iems.nwu.edu/~nelsonb/ARTA/>.

Figure 6: Time-series Plot of Observations from the ARTA-process Model

Figure 8: Scatterplot of (Y_t, Y_{t+1}) for the ARTA-process Model

Figure 7: Scatterplot of (Y_t, Y_{t+1}) for the Empirical Data

ACKNOWLEDGMENTS

The author would like to thank Barry L. Nelson of Northwestern University for his active participation in this research and in writing this paper.

REFERENCES

- Cario, M. C., and B. L. Nelson. 1996. Autoregressive to anything: Time-series input processes for simulation. *Operations Research Letters*, in press.
- Lewis, P. A. W., E. McKenzie, and D. K. Hugus. 1989. Gamma processes. *Communications in Statistics-Stochastic Models* 5:1-30.
- Li, S. T., and J. L. Hammond. 1975. Generation of pseudorandom numbers with specified univariate distributions and correlation coefficients. *IEEE*

Transactions on Systems, Man and Cybernetics 5:557-561.

- Livny, M., B. Melamed, and A. K. Tsolis. 1993. The impact of autocorrelation on queuing systems. *Management Science* 39:322-339.
- Melamed, B. 1991. TES: A class of methods for generating autocorrelated uniform variates. *ORSA Journal on Computing* 3:317-329.
- Melamed, B., J. R. Hill, and D. Goldsman. 1992. The TES methodology: Modeling empirical stationary time series. In *Proceedings of the 1992 Winter Simulation Conference*, ed. J. J. Swain, D. Goldsman, R. C. Crain and J. R. Wilson, 135-144. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey.
- Piessens, R., E. de Doncker-Kapenga, C. W. Überhuber, and D. K. Kahaner. 1983. *Quadpack: A subroutine package for automatic integration*. New York: Springer-Verlag.
- Song, W. T., L. Hsiao, and Y. Chen. 1996. Generation of pseudo-random time series with specified marginal distributions. *European Journal of Operational Research*, forthcoming.
- Whitt, W. 1976. Bivariate distributions with given marginals. *The Annals of Statistics* 4:1280-1289.

AUTHOR BIOGRAPHY

MARNE C. CARIO is a systems analyst at Delphi Packard Electric Systems in Warren, Ohio. Her research interests focus on simulation input modeling and the application of simulation to manufacturing systems.