

A FLEXIBLE AUTOMATED PROCEDURE FOR MODELING COMPLEX ARRIVAL PROCESSES

Michael E. Kuhl
Sachin G. Sumant

Dept. of Industrial & Systems Engineering
81 Lomb Memorial Drive
Rochester Institute of Technology
Rochester, NY 14623, U.S.A.

James R. Wilson

Dept. of Industrial Engineering
2401 Stinson Drive, Riddick Labs 328
North Carolina State University
Raleigh, NC 27695-7906, U.S.A.

ABSTRACT

To automate the multiresolution procedure of Kuhl and Wilson for modeling and simulating arrival processes that exhibit long-term trends and nested periodic effects (such as daily, weekly, and monthly cycles), we present a statistical-estimation method that involves the following steps at each resolution level corresponding to a basic cycle: (a) transforming the cumulative relative frequency of arrivals within the cycle (for example, the percentage of all arrivals as a function of the day of the week within the weekly cycle) to obtain a statistical model with normal, constant-variance responses; (b) fitting a specially formulated polynomial to the transformed responses; (c) performing a likelihood ratio test to determine the degree of the fitted polynomial; and (d) fitting a polynomial of the degree determined in (c) to the original (untransformed) responses. An example demonstrates web-based software that implements this flexible approach to handling complex arrival processes.

1 INTRODUCTION

Time-varying arrival processes are routinely encountered in practical applications of industrial and systems engineering techniques. The following are typical situations in which the arrival rate of relevant entities depends strongly on time: demands for seasonal products such as lawn mowers; arrivals of patrons at an amusement park; arrivals of patients at an emergency room; and the arrivals of telephone calls at a customer service center. To analyze or improve system operation in such situations, discrete-event stochastic simulation is often the technique of choice. Consequently, high-fidelity probabilistic input models are often needed to perform meaningful simulation experiments. In the past, nonhomogeneous Poisson processes (NHPPs) have been used successfully to model complex time-dependent arrival processes (Kuhl, Damerdj, and Wilson 1997; Kuhl and Wilson 2000, 2001).

An NHPP $\{N(t): t \geq 0\}$ is a counting process such that $N(t)$ is the number of arrivals in the time interval $(0, t]$; and $\lambda(t)$, the instantaneous arrival rate at time t , is a nonnegative function of time satisfying the usual Poisson postulates so that the corresponding (cumulative) mean value function is given by

$$\mu(t) \equiv E[N(t)] = \int_0^t \lambda(z) dz \quad \text{for all } t \geq 0.$$

The rate or mean value function of the NHPP $\{N(t), t \geq 0\}$ completely characterizes the probabilistic behavior of the process (Çınlar 1975).

Both parametric and nonparametric methods have been developed to estimate the rate or mean value function of the process $\{N(t) : t \geq 0\}$ from observed data. This paper focuses on the nonparametric “multiresolution” estimation method of Kuhl, Damerdj, and Wilson (1997) and Kuhl and Wilson (2001), where the target arrival process may exhibit a long-term trend as well as nested cyclic effects that do not necessarily possess the symmetry of sinusoidal oscillations. For example, in developing a large-scale simulation model of the organ-procurement and patient-registration processes for liver transplants in the United States, Pritsker et. al (1995) found that the arrival streams of liver donors and liver patients have pronounced long-term trends as well as asymmetric cyclic effects that depend on the time of year, the day of the week, and the hour of the day; see also Pritsker (1998).

In the multiresolution procedure of Kuhl, Damerdj, and Wilson (1997) and Kuhl and Wilson (2001), we begin at Resolution 0 by estimating any long-term trend that is present in the overall observation interval $[0, S]$. If p cyclic effects are also present, then we evaluate the arrival process on the scale of the largest cycle (that is, the cycle of length b_1 which is associated with Resolution 1) to estimate the major periodic features of the arrival process; and we assume S/b_1 is an integer so that the overall observation interval consists of complete Resolution-1 cycles.

Then we evaluate the effects corresponding to cycles of progressively smaller lengths that are nested within each other so that at Resolution i (for $i = 1, 2, \dots, p - 1$), the associated cycle length b_i consists of an integral number of cycles of length b_{i+1} which are associated with Resolution $i + 1$. Thus the cycle length b_p at Resolution p is the smallest observation subinterval over which we seek to model accurately the periodic behavior of the target arrival process. Finally we construct the estimated mean value function from the estimates of all the cyclic effects as well as the long-term trend effect.

Kuhl and Wilson (2001) develop a theoretical basis for the multiresolution procedure, but they do not provide a specific implementation of the procedure that can be readily applied to sample data. In this paper, we develop a completely automated implementation of the multiresolution procedure that can be executed over the web, enabling users to fit their arrival processes and then generate an independent realization of the fitted process interactively without the need for downloading and installing the multiresolution software.

The rest of this paper is organized as follows. For completeness, in §2 we summarize the mathematical formulation of the multiresolution procedure. In §3 we present the automated fitting procedure, which involves the following steps at each resolution level associated with a basic cycle:

- (a) making a variance stabilizing transformation (specifically, the composite of the arc-sine and square-root transformations) of the cumulative relative frequency of arrivals within the cycle so as to obtain normally distributed responses with a constant variance;
- (b) fitting to the transformed arrival data a polynomial that has been specially formulated to ensure it has the properties of a cumulative distribution function (c.d.f.);
- (c) performing a likelihood ratio test to determine the appropriate degree of the fitted polynomial in step (b); and
- (d) fitting to the original (untransformed) arrival data a polynomial of the degree determined in step (c).

In §4 we present an example that illustrates the automated multiresolution procedure and gives some indication of the flexibility of this method. In §5 we demonstrate our web-based input modeling software by applying it to the example detailed in §4. Finally conclusions and recommendations for future work are summarized in §6.

2 SUMMARY OF THE MULTIRESOLUTION ESTIMATION PROCEDURE

The mean value function is fit to data observed over the time interval $[0, S]$, where the data may contain a long-

term trend or p cyclic effects having periods of length b_i ($i = 1, 2, \dots, p$), respectively. This method can be applied to arrival processes satisfying two key assumptions that are formally stated below. As an aid in explaining the significance of these assumptions, we consider a simple example of an arrival process whose instantaneous arrival rate has a long-term trend over a time horizon of 28 days as well as weekly and daily cyclic effects—that is, periodic rate components with periods of 1 week and 1 day.

Assumption 1. There are p distinct cycle lengths (periods) $b_1 > b_2 > \dots > b_p$ such that b_i is an integral multiple of b_{i+1} for $i = 1, 2, \dots, p - 1$. Moreover, we take the time horizon $[0, S]$ such that S is an integral multiple of b_1 ; and we let $b_0 \equiv S$.

In our example, we take 1 day as the unit of time so that we have an overall time horizon of $S = b_0 = 28$ days; and within each (weekly) cycle of length $b_1 = 7$ days, there are (daily) cycles of length $b_2 = 1$ day so that Assumption 1 is satisfied. Notice that Assumption 1 requires knowledge of the cycle lengths of all the cyclic effects observed in the arrival process; and in our experience, such information is often available in practical applications.

Assumption 2. Within the j th cycle $[(j-1)b_i, jb_i]$ of length b_i (for $i = 1, 2, \dots, p$), the arrival rate at time t is proportional to a single baseline function $\lambda_i(s)$, where $s = t - (j-1)b_i$ is the offset from the beginning of the cycle so that we have $\lambda(t) = \alpha_{i,j} \lambda_i(t - (j-1)b_i) = \alpha_{i,j} \lambda_i(s)$, and $\{\alpha_{i,\ell} : \ell = 1, 2, \dots, S/b_i\}$ are the constants of proportionality for all cycles of length b_i .

In our example, Assumption 2 requires that the rate of arrivals within each day should follow the same general profile as a function of the time of day; but in different days within each weekly cycle, the daily arrival rates may have different scaling factors $\{\alpha_{2,\ell} : \ell = 1, 2, \dots, b_1/b_2 = 7\}$ to account for different numbers of expected arrivals on different days of the week. Similarly, Assumption 2 requires that the rate of arrivals within each week should follow the same general profile as a function of the time of the week; but in different weeks within the overall time horizon of 28 days, the weekly arrival rates may have different scaling factors $\{\alpha_{1,\ell} : \ell = 1, 2, \dots, b_0/b_1 = 4\}$ to account for different numbers of expected arrivals on different weeks in the time horizon. In §4 we elaborate this example to illustrate not only the probabilistic structure of NHPPs satisfying Assumptions 1 and 2 but also the general approach we take to modeling and estimation of such arrival processes.

It is easy to see that Assumption 2 holds if and only if the target arrival process has the following property: within each cycle $[(j-1)b_i, jb_i]$ of length b_i ($i = 1, \dots, p$), the buildup to time $t = (j-1)b_i + s$ (where $s = t \bmod b_i$) of the cumulative fraction $[\mu(t) - \mu((j-1)b_i)]/[\mu(jb_i) - \mu((j-1)b_i)]$ of the expected number of arrivals in the cycle is described

by the same function $R_i(s)$ of the time s that has elapsed since the beginning of the cycle. Thus to estimate the mean value function $\mu(t)$, we first need to estimate the function

$$R_i(s) \equiv \frac{\int_0^s \lambda_i(z) dz}{\int_0^{b_i} \lambda_i(z) dz} \quad (1)$$

for $s \in [0, b_i]$ and $i = 0, 1, 2, \dots, p$.

At each resolution level i , we must obtain an estimator $\hat{R}_i(s)$ of the function $R_i(s)$ with the following properties: the initial value $\hat{R}_i(0) = R_i(0) = 0$; the final value $\hat{R}_i(b_i) = R_i(b_i) = 1$; and the derivative $\hat{R}'_i(s) > 0$ for $s \in [0, b_i]$ and $i = 0, 1, 2, \dots, p$. At resolution level 0, we fit a monotonically increasing function $\hat{R}_0(t)$, $t \in [0, S]$, to the points $\left\{ [jb_1, N(jb_1)/N(S)]^T : j = 0, 1, \dots, S/b_1 \right\}$.

At resolution level i for $i = 1, 2, \dots, p-1$, we construct a monotonically increasing function $\hat{R}_i(s)$ to estimate the cumulative percentage of arrivals that are expected to occur during the first s time units of the associated cycle of length b_i , where $s \in [0, b_i]$. That is, for resolution level i we fit the function $\hat{R}_i(s)$ to the points $\left\{ [jb_{i+1}, G_i(jb_{i+1})]^T : j = 0, 1, \dots, b_i/b_{i+1} \right\}$, where

$$G_i(s) = \frac{1}{N(S)} \sum_{\ell=0}^{(S/b_i)-1} [N(\ell b_i + s) - N(\ell b_i)] \quad (2)$$

for all $s \in [0, b_i]$.

At resolution p , we let $\{\tau_k : k = 1, 2, \dots, N(S)\}$ denote the observed arrival times; and we define $\eta_k = \tau_k \bmod b_p$ for $k = 1, \dots, N(S)$. Let $\eta_{(1)} \leq \dots \leq \eta_{(k)}$ denote the ordered arrival times within the level- p cycle $[0, b_p]$. We then fit a monotonically increasing function $\hat{R}_p(s)$ to the points $\left\{ [\eta_{(k)}, k/N(S)]^T : k = 1, \dots, N(S) \right\}$ subject to the usual boundary conditions that $\hat{R}_p(0) = 0$ and $\hat{R}_p(b_p) = 1$. The final estimate $\hat{\mu}(t)$ of the mean value function is computed as follows:

$$\hat{\mu}(t) = N(S) \hat{Q}_0(t) \text{ for } t \in [0, S],$$

where the functions $\left\{ \hat{Q}_i(t) : i = p, p-1, \dots, 1, 0 \right\}$ are defined iteratively by

$$\hat{Q}_p(t) = \hat{R}_p(t - (j_{p,t} - 1)b_p),$$

and for $i = p-1, \dots, 1, 0$, we take

$$\begin{aligned} \hat{Q}_i(t) = & \hat{R}_i((j_{i+1,t} - 1)b_{i+1} - (j_{i,t} - 1)b_i) \\ & + \left[\hat{R}_i(j_{i+1,t}b_{i+1} - (j_{i,t} - 1)b_i) \right. \\ & \left. - \hat{R}_i((j_{i+1,t} - 1)b_{i+1} - (j_{i,t} - 1)b_i) \right] \hat{Q}_{i+1}(t), \end{aligned}$$

where the interval containing t is within each resolution is given by the index,

$$j_{i,t} \text{ is the unique integer } j \text{ such that } (j-1)b_i \leq t < jb_i.$$

Although the multiresolution method proposed by Kuhl and Wilson (2001) does not require a specific functional form for the estimator $\hat{R}_i(s)$ of the function $R_i(s)$ at each level of resolution, in the next section we specify a complete methodology for implementing the multiresolution procedure in general applications.

3 AUTOMATED IMPLEMENTATION OF THE MULTIREOLUTION PROCEDURE

3.1 Selecting the Functional Form for Each $\hat{R}_i(\cdot)$

The multiresolution procedure requires that each fitted function $\hat{R}_i(s)$ is monotone and that $\hat{R}_i(0) = 0$ and $\hat{R}_i(b_i) = 1$ for $i = 0, 1, \dots, p$. Therefore, we utilized a special form of a polynomial that meets these requirements—namely, a degree r polynomial of the form

$$\hat{R}_i(s) \equiv \begin{cases} s/b_i, & \text{if } r = 1, \\ \sum_{k=1}^{r-1} \beta_k (s/b_i)^k + \left(1 - \sum_{k=1}^{r-1} \beta_k \right) (s/b_i)^r & \text{if } r > 1, \end{cases} \quad (3)$$

for $s \in [0, b_i]$ and $i = 0, 1, 2, \dots, p$. We constrain the coefficients $\{\beta_k : k = 1, \dots, r-1\}$ in (3) to yield $\hat{R}'_i(s) > 0$ for all $s \in [0, b_i]$. This polynomial automatically satisfies the required boundary conditions. Further, the simplest process we can model is the one with a constant arrival rate over the interval $[0, b_i]$, which corresponds to the simplest form of our model (3) wherein we take $r = 1$ (that is, a linear mean value function).

To estimate the function (1) by a fitted function $\hat{R}_i(s)$ of the form (3) at each resolution level i , we must determine the appropriate degree r of the polynomial (3) and then estimate the associated polynomial coefficients $\beta_1, \beta_2, \dots, \beta_{r-1}$. The standard approach for solving such a problem is to perform regression analysis using a forward selection or backward elimination procedure (Draper and Smith 1998).

In a forward selection procedure, we fit polynomials of degree $r - 1$ and r to the sample data; then we compute the mean square error for each fit and perform an appropriate likelihood ratio test to determine if significant improvement in the fit is achieved by increasing the degree of the fitted polynomial from $r - 1$ to r . If a significantly better fit is obtained with the degree- r polynomial, then we continue the fitting procedure by incrementing the value of r and repeating the likelihood ratio test; otherwise, we terminate the fitting procedure, delivering $r - 1$ as the degree of the fitted polynomial.

However, the assumptions of classical regression analysis require observations of a dependent variable that are independent and normal with a constant variance so that the corresponding residuals are independent and identically distributed normal random variables. In the case of fitting arrival data obtained from an NHPP that satisfies Assumptions 1 and 2, the observed cumulative relative frequencies within each basic cycle are neither normal nor independent; moreover, such responses do not possess a constant variance. Therefore, a variance stabilizing transformation must be applied to the data before we try to use standard statistical procedures to determine an appropriate degree r for the function $\hat{R}_i(s)$.

3.2 Likelihood Ratio Test for Determining the Degree of Each $\hat{R}_i(s)$

At each resolution level i , we let $m_i \equiv b_i/b_{i+1} - 1$ for $0 \leq i \leq p - 1$ and $m_i \equiv N(S)$ for $i = p$; and we seek to fit a polynomial function $\hat{R}_i(s)$ of the form (3) with appropriate degree r to a set of points having the general form

$$\{(Z_j, W_j)^T : j=1, \dots, m_i\} \quad (4)$$

as defined in §2. For example if the resolution level i is in the range $\{1, \dots, p - 1\}$, then at the j th point in (4) we take the abscissa $Z_j = jb_{i+1}$ and the ordinate $W_j = G_i(jb_{i+1})$ for $j = 1, \dots, m_i$. Since W_j is always a proportion, we exploit the variance stabilizing transformation

$$Y_j \equiv \sin^{-1} \left[\sqrt{W_j} \right] \quad \text{for } j=1, \dots, m_i; \quad (5)$$

see pp. 231–238 of Box, Hunter, and Hunter (1978) and pp. 291–294 of Draper and Smith (1998).

Corresponding to the dependent variable Y_j defined by (5), we define the independent variable

$$X_j \equiv Z_j / b_i = jb_{i+1} / b_i \quad \text{for } j = 1, \dots, m_i; \quad (6)$$

and in terms of the vector of regression coefficients

$$\mathbf{C}_r \equiv \begin{cases} \pi/2, & \text{if } r = 1, \\ [C_1, C_2, \dots, C_{r-1}], & \text{if } r \geq 2, \end{cases} \quad (7)$$

for the degree- r polynomial

$$\Gamma_r(u; \mathbf{C}_r) \equiv \begin{cases} \left(\frac{\pi}{2}\right)u, & \text{if } r=1, \\ \sum_{k=1}^{r-1} C_k u^k + \left(\frac{\pi}{2} - \sum_{k=1}^{r-1} C_k\right)u^r, & \text{if } r>1, \end{cases} \quad (8)$$

we postulate the following statistical model for Y_j as a function of X_j ,

$$Y_j = \Gamma_r(X_j; \mathbf{C}_r) + \varepsilon_j \quad \text{for } j = 1, 2, \dots, m_i. \quad (9)$$

If (9) is valid, then the transformation (5) yields approximately normal residuals with constant variance σ^2 so that

$$\{\varepsilon_j : j=1, \dots, m_i\} \sim N(0, \sigma^2). \quad (10)$$

Notice that in (8) we have

$$\Gamma_r(0; \mathbf{C}_r) = 0 \quad \text{and} \quad \Gamma_r(1; \mathbf{C}_r) = \pi/2. \quad (11)$$

For fixed r ($r = 2, 3, \dots$), the estimator $\tilde{\mathbf{C}}_r$ of the coefficient vector \mathbf{C}_r is the solution of the following constrained least-squares regression problem:

$$\min_{\hat{\mathbf{C}}_r} \sum_{j=1}^{m_i} [Y_j - \Gamma_r(X_j; \hat{\mathbf{C}}_r)]^2 \quad (12)$$

subject to

$$\Gamma'_r(u; \hat{\mathbf{C}}_r) \neq 0 \quad \text{for all } u \in (0, 1). \quad (13)$$

Observe that

$$\Gamma'_r(u; \hat{\mathbf{C}}_r) = \sum_{k=1}^{r-1} k \hat{C}_k u^{k-1} + r \left(\frac{\pi}{2} - \sum_{k=1}^{r-1} \hat{C}_k \right) u^{r-1}, \quad (14)$$

so that (13) can be interpreted as requiring the zeros of the degree- $(r - 1)$ polynomial (14) to lie outside the unit interval.

Associated with (12) is the error sum of squares

$$\text{SSE}_r = \sum_{j=1}^{m_i} [Y_j - \Gamma_r(X_j; \tilde{\mathbf{C}}_r)]^2 \quad (15)$$

for $r = 1, 2, \dots$, where we take $\tilde{\mathbf{C}}_1 \equiv (\pi/2)$ in conformance with (7). Now for fixed r ($r \geq 2$), $\tilde{\mathbf{C}}_r$ is the maximum likelihood estimator of \mathbf{C}_r when (9)–(10) hold and the residuals are independent; moreover, we see that

$$\tilde{\sigma}_r^2 \equiv \text{SSE}_r / m_i \quad (16)$$

is the maximum likelihood estimator of σ^2 (Arnold 1981). Given $\mathbf{Y} \equiv [Y_1, \dots, Y_{m_i}]^T = \mathbf{y} \equiv [y_1, \dots, y_{m_i}]^T$, we see that the associated likelihood function is given by

$$L_r(\tilde{\mathbf{C}}_r; \mathbf{y}) = \prod_{j=1}^{m_i} \frac{1}{\tilde{\sigma}_r \sqrt{2\pi}} \exp \left\{ -\frac{1}{2} \left[\frac{Y_j - \Gamma_r(X_j; \tilde{\mathbf{C}}_r)}{\tilde{\sigma}_r} \right]^2 \right\} \quad (17)$$

$$= (2\pi)^{-m_i/2} (\tilde{\sigma}_r^2)^{-m_i/2} \exp \left[-\frac{1}{2} \left(\frac{\text{SSE}_r}{\tilde{\sigma}_r^2} \right) \right]$$

$$= (2\pi e \tilde{\sigma}_r^2)^{-m_i/2} \quad (18)$$

so that the resulting log-likelihood function for the fixed degree r is

$$\mathcal{L}_r(\tilde{\mathbf{C}}_r; \mathbf{y}) = -\frac{m_i}{2} \left[\ln(2\pi) + 1 + \ln(\tilde{\sigma}_r^2) \right]. \quad (19)$$

The degree r of the polynomial (8) is determined by a likelihood ratio test that has been adapted to constrained nonlinear regression. This approach to determining r is based on a similar technique used by Avramidis and Wilson (1994) in the context of another constrained nonlinear regression problem. At the outset, we assume that (8)–(10) hold for some value of r to be determined. Starting with the degree $r = 2$ and computing the optimal solution $\tilde{\mathbf{C}}_r$ to (12), we seek to test the null hypothesis that

$$\sum_{k=1}^{r-1} C_k = \frac{\pi}{2} \quad (20)$$

(so that the degree of the polynomial (8) is at most $r - 1$) versus the alternative hypothesis that

$$\sum_{k=1}^{r-1} C_k \neq \frac{\pi}{2} \quad (21)$$

(so that the degree of the polynomial (8) is at least r). If (20) holds, then Theorem 4.4.4 of Serfling (1980) ensures that

$$2[\mathcal{L}_r(\tilde{\mathbf{C}}_r; \mathbf{y}) - \mathcal{L}_{r-1}(\tilde{\mathbf{C}}_{r-1}; \mathbf{y})] \xrightarrow[m_i \rightarrow \infty]{D} \chi^2(1), \quad (22)$$

where $\chi^2(1)$ is a chi-square random variable with 1 degree of freedom. In view of (19), we see that

$$2[\mathcal{L}_r(\tilde{\mathbf{C}}_r; \mathbf{y}) - \mathcal{L}_{r-1}(\tilde{\mathbf{C}}_{r-1}; \mathbf{y})] = -m_i \ln(\tilde{\sigma}_r^2 / \tilde{\sigma}_{r-1}^2); \quad (23)$$

and in the formal algorithmic statement of our likelihood ratio procedure for estimating r given below, we also take

$$\text{SSE}_1 \equiv \sum_{j=1}^{m_i} \left[Y_j - \frac{\pi j}{2(m_i + 1)} \right]^2 \quad (24)$$

and

$$\tilde{\sigma}_1^2 \equiv \frac{\text{SSE}_1}{m_i}. \quad (25)$$

We perform the likelihood ratio test at the level of significance α , where $0 < \alpha < 1$, and the final estimate of r is

$$\tilde{r} = \min \left\{ r : r \geq 2 \text{ and } -m_i \ln \left(\frac{\tilde{\sigma}_r^2}{\tilde{\sigma}_{r-1}^2} \right) \leq \chi_{1-\alpha}^2(1) \right\} - 1, \quad (26)$$

where $\chi_{1-\alpha}^2(1)$ is the $1 - \alpha$ quantile of the chi-square distribution with 1 degree of freedom. After determining the degree \tilde{r} of the polynomial (8) used to fit the arrival data that was transformed via (5), we take the same degree \tilde{r} for the polynomial (3) used to fit the original (untransformed) arrival data defined in (4).

Kuhl and Wilson (2000) formulated and evaluated an ordinary least squares (OLS) procedure for estimating the parametric mean value function of a NHPP and demonstrated that the estimation accuracy and computational efficiency of the OLS procedure in this type of application is reasonable. The same OLS procedure is applied here. The likelihood ratio test algorithm is summarized in Figure 1.

4 ILLUSTRATIVE EXAMPLE

To illustrate the automated multiresolution procedure for estimating NHPPs, we elaborate the example introduced in §2. Recall that, the arrival rate has a long-term trend over a time horizon of 28 days as well as weekly and daily cyclic effects. Since the arrival process has two periodic components, we estimated the function (1) at the following levels of detail: (a) Resolution 0, corresponding to the

Likelihood Ratio Test Algorithm

- [1] Take the composite (5) of the arc-sine and square-root transformations of the original arrival data (4) within each cycle.
- [2] Set $r \leftarrow 1$ and check the maximum value of \tilde{r}
If $m_i \leq 2$, then deliver $\tilde{r} \leftarrow 1$ and stop; else set $r \leftarrow 2$ and go to [3].
- [3] Compute OLS estimator \tilde{C}_r .
a) Use starting value $\hat{C}_r = (\tilde{C}_{r-1}, 0)$ to obtain \tilde{C}_r
(e.g., to compute \tilde{C}_2 , start with $\hat{C}_2 = (\pi/2, 0)$).
b) Determine the estimator \tilde{C}_r that minimizes SSE_r as given by (15) subject to (13).
- [4] Calculate $\tilde{\sigma}_{r-1}^2 = SSE_{r-1}/m_i$ and $\tilde{\sigma}_r^2 = SSE_r/m_i$.
- [5] If $-m_i \ln \left(\frac{\tilde{\sigma}_r^2}{\tilde{\sigma}_{r-1}^2} \right) \leq \chi_{1-\alpha}^2(1)$, then deliver $\tilde{r} \leftarrow r-1$
and stop; otherwise, set $r \leftarrow r+1$ and go to [3].

Figure 1: Likelihood Ratio Test Algorithm

long-term trend over a period of 28 days; (b) Resolution 1, corresponding to the cyclic rate component with a period of 1 week; and (c) Resolution 2, corresponding to the cyclic rate component with a period of 1 day. The input data file contained 216 arrival times, and we specified the significance level $\alpha=0.10$ for the likelihood ratio test (26).

At Resolution 0, the number of points involved in the fitting procedure is $m_0=3$, excluding both endpoints. These 3 points correspond to the cumulative fraction of arrivals observed at the end of each week. Applying the likelihood ratio test of Figure 1, we see that $\hat{\Gamma}_0(s)$, the fitted function at resolution 0 for the transformed data, is a linear function (that is, $\tilde{r}_0 = 1$). Consequently, a degree-1 polynomial $\hat{R}_0(s)$ is fit to the original data for the overall observation interval. Figures 2 and 5 respectively show the fitted functions $\hat{\Gamma}_0(s)$ and $\hat{R}_0(s)$ at Resolution 0.

Next, we estimate the Resolution-1 function $R_1(s)$ associated with a period of 1 week. First the arrivals over each week were superimposed. There are $m_1=6$ points corresponding to the cumulative fraction of arrivals observed at the end of each day. These data are then transformed and the fitting procedure is applied, resulting in a cubic function being fit to the data so that $\tilde{r}_1=3$. The fitted function at Resolution 1 for the transformed data, $\hat{\Gamma}_1(s)$, is shown in Figure 3. A corresponding degree-3 polynomial $\hat{R}_1(s)$ is then fit to the original data, as shown in Figure 6.

The next step in the procedure is to estimate the resolution-2 function $R_2(s)$ associated with a period of 1 day. The $m_2=216$ data points corresponding to the cumulative fraction of arrivals are superimposed over each day. Thus at Resolution 2, the fitted function $\hat{\Gamma}_2(s)$ is a polynomial of degree $\tilde{r}_2=10$ as shown in Figure 4. The corresponding estimated function $\hat{R}_2(s)$ for the original data at Resolution 2 is shown in Figure 7.

The estimates of the functions for the transformed data are:

Resolution 0:

$$\hat{\Gamma}_0(s) = 0.0560998693s \text{ for all } s \in [0, 28]$$

Resolution 1:

$$\hat{\Gamma}_1(s) = 0.552856266s - 0.102613039s^2 + 0.00795580633s^3 \text{ for } s \in [0, 7]$$

Resolution 2:

$$\begin{aligned} \hat{\Gamma}_2(s) = & 7.58305597s - 40.6680336s^2 + 124.117729s^3 \\ & - 194.861145s^4 + 137.300613s^5 - 6.46647882s^6 \\ & - 25.725769s^7 - 18.9463806s^8 + 20.8811493s^9 \\ & - 1.64394438s^{10} \text{ for } s \in [0, 1]. \end{aligned}$$

The estimates of the functions for the original data are:

Resolution 0:

$$\hat{R}_0(s) = 0.0357142857s \text{ for } s \in [0, 28]$$

Resolution 1:

$$\hat{R}_1(s) = 0.264645542s - 0.0184986202s^2 + 0.0001571824s^3 \text{ for } s \in [0, 7]$$

Resolution 2:

$$\begin{aligned} \hat{R}_2(s) = & 2.54902697s - 8.86513996s^2 + 28.077919s^3 \\ & + 46.5554962s^4 + 36.8081818s^5 - 11.000386s^6 \\ & - 2.07357478s^7 + 3.56143045s^8 - 1.48909688s^9 \\ & - 0.0128641129s^{10} \text{ for } s \in [0, 1]. \end{aligned}$$

After fitting the required functions at each resolution, we constructed the estimated mean value function. The fitted mean value function is plotted against the observed cumulative number of arrivals over the entire time horizon $[0, 28]$ as shown in Figure 8. We see that the multiresolution procedure is capable of accurately estimating the mean value function as well as the auxiliary functions $\{R_i(s): i=1, \dots, p\}$ from one realization of the process.

The estimation procedure illustrated in this paper is controllable in the sense the user can specify the maximum degree of the polynomial fit to the data at each resolution, and the user can specify the significance level for the likelihood ratio test. The above example with three resolutions is a fairly complex application which clearly illustrates the flexibility of the fitting procedure.




Figure 2: Resolution 0, Estimated Function for Transformed Data




Figure 5: Resolution 0, Estimated Function for Original (Untransformed) Data




Figure 3: Resolution 1, Estimated Function for Transformed Data




Figure 6: Resolution 1, Estimated Function for Original (Untransformed) Data




Figure 4: Resolution 2, Estimated Function for Transformed Data




Figure 7: Resolution 2, Estimated Function for Original (Untransformed) Data




Figure 8: Fitted Mean Value Function

At Resolution 0, Figure 5 shows a reasonable linear fit to the original (untransformed) data. Similarly, at Resolution 1, Figure 6 reveals that the fitted cubic function provides an excellent fit to the original data. The most difficult task was to obtain an adequate approximation to the 216 arrival times at Resolution 2. Figure 7 shows that degree-10 polynomial obtained via the estimation method of Figure 1 provides an excellent approximation to the distribution of arrival times during the day. Finally the fitted mean value function in Figure 8 demonstrates that the overall arrival process is accurately modeled and that the likelihood ratio test performed well in this application.

Note that the likelihood ratio test is only be applied to determine the degree of the polynomial approximation $\hat{R}_i(s)$ to the function $R_i(s)$ at each resolution level i . It is not designed to perform any other task in the multiresolution procedure.

5 WEB-BASED INPUT MODELING

The multiresolution procedure described in the previous sections has been implemented in a web-based environment that is accessible via <http://www.rit.edu/~kuhl1/simulation>. The web site can be used to fit an NHPP to data using the multiresolution procedure as well as to generate arrival times from the NHPP that could be used as inputs to a simulation model.

The benefits of the web site include the ability for a simulation analyst to access the software for fitting NHPPs to data from anywhere in the world. In addition, the data analysis can be completed through the use of a web browser without the need to install the fitting software on individual computers. Furthermore, the web site can be used in simulation courses to teach students about NHPPs and data fitting without having to provide students with their own copy of the fitting software.

Figure 9 illustrates a sample screen obtained by visiting the web site.

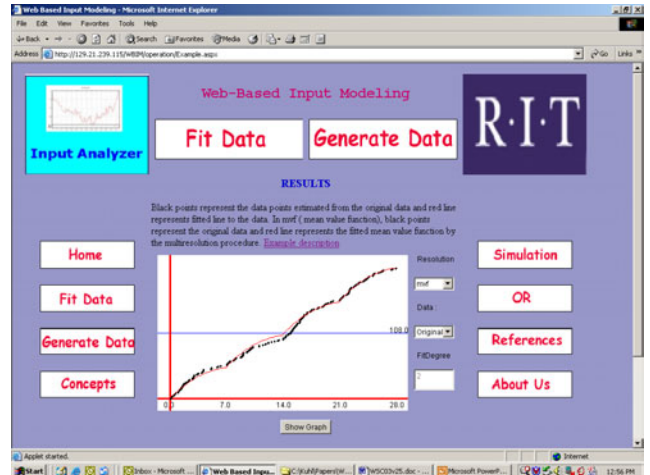


Figure 9: Web Based Input Modeling Software

To utilize the fitting procedure, the user will need to supply an ASCII (plaintext) file containing a single column of observed arrival times. The user will also be prompted for the following information about the data and the parameters for the fitting procedure:

- The ending time of the observation interval;
- The number of periodic (cyclic) components;
- The period of each cyclic component;
- The significance level α that will be used in the likelihood ratio test (26) for determining the degree of the polynomial; and
- The maximum degree of the polynomials to fit at each resolution.

Once the user has submitted this information, the fitting procedure is run and the results of the fit are reported including the fitted polynomial at each resolution, graphical results of the fitted data at each resolution for both the transformed and original data sets, and a graph of the fitted mean value function versus the observed cumulative arrivals.

After fitting an NHPP to arrival data, we can use the web site to generate a realization of the fitted arrival process which can in turn drive a simulation experiment. The resulting arrival times are written into an ASCII (plaintext) file in a single-column format; and the resulting file can be easily read into commercial simulation software packages during the simulation run.

6 CONCLUSION AND RECOMMENDATIONS

In this paper we have developed a completely automated procedure for estimating arrival processes that exhibit long-term trends and nested periodic effects. We have also provided some evidence that the procedure is sufficiently flexible to enable accurate estimation of the mean value function of the target arrival process while using a relatively small number of parameters that can be efficiently estimated from the history of process.

Future extensions of the multiresolution procedure will include (a) a graphical user interface for invoking the procedure in popular simulation software systems; and (b) efficient software for generating independent replications of a fitted arrival process. Another direction for future work is to develop statistical tools for validating the basic assumptions of the multiresolution procedure in arrival processes to which the procedure might be applied.

REFERENCES

- Arnold, S. F. 1981. *The theory of linear models and multivariate analysis*. New York: Wiley.
- Avramidis, A. N., and J. R. Wilson. 1994. A flexible method for estimating inverse distribution functions in simulation experiments. *ORSA Journal on Computing* 6 (4): 342–355.
- Box, G. E. P., W. G. Hunter, and J. S. Hunter. 1978. *Statistics for experimenters: An introduction to design, data analysis, and model building*. New York: Wiley.
- Çınlar, E. 1975. *Introduction to stochastic processes*. New Jersey: Prentice-Hall.
- Draper, N. R., and H. Smith. 1998. *Applied regression analysis*. 3d ed. New York: Wiley.
- Kuhl, M. E., H. Damerdjji, and J. R. Wilson. 1997. Estimating and simulating Poisson processes with trends or asymmetric cyclic effects. In *Proceedings of the 1997 Winter Simulation Conference*, ed. S. Andradóttir, K. J. Healy, D. H. Withers, and B. L. Nelson, 287–295. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers. Available online via www.informs-cs.org/wsc97papers/0287.PDF.
- Kuhl, M. E., and J. R. Wilson. 2000. Least squares estimation of nonhomogeneous Poisson processes. *Journal of Statistical Computation and Simulation* 67:75–108.
- Kuhl, M. E., and J. R. Wilson. 2001. Modeling and simulating Poisson processes having trends or nontrigonometric cyclic effects. *European Journal of Operational Research* 133 (3): 566–582.
- Pritsker, A. A. B., D. L. Martin, J. S. Reust, M. A. F. Wagner, J. R. Wilson, M. E. Kuhl, J. P. Roberts, O. P. Daily, A. M. Harper, E. B. Edwards, L. Bennett, J. F. Burdick, and M. D. Allen. 1995. Organ transplantation policy evaluation. In *Proceedings of the 1995 Winter Simulation Conference*, ed. C. Alexopoulos, K. Kang, W. Lilegdon, and D. Goldsman, 1314–1323. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.
- Pritsker, A. A. B. 1998. Life and death decisions: Organ transplantation allocation policy analysis. *OR/MS Today* 25:22–28.
- Serfling, R. J. 1980. *Approximation theorems of mathematical statistics*. New York: Wiley.

AUTHOR BIOGRAPHIES

MICHAEL E. KUHL is an assistant professor in the Department of Industrial and Systems Engineering and an assistant professor of Microsystems Engineering at Rochester Institute of Technology. His research interests include simulation input modeling and simulation applied to semiconductor/microsystems fabrication and health care systems. He is a member of IIE, INFORMS, and ASEE. His e-mail address is mekeie@rit.edu, and his web page is www.rit.edu/~mekeie/.

SACHIN G. SUMANT is a master of science student in the Department of Industrial and Systems Engineering at Rochester Institute of Technology. His research interests include simulation input modeling and web-based simulation. He is a member of INFORMS and ASME. His e-mail address is sumantsachin@hotmail.com.

JAMES R. WILSON is Professor and Head of the Industrial Engineering Department at North Carolina State University. His research interests are focused on probabilistic and statistical issues in the design and analysis of simulation experiments. Currently he serves as chair of the WSC Board of Directors, corepresenting INFORMS–CS. He is a member of ASA and INFORMS, and he is a Fellow of IIE. His e-mail address is jwilson@eos.ncsu.edu, and his web page is www.ie.ncsu.edu/jwilson.