5 Models of Variability

Modeling always involves some degree of approximation. We have seen how complex and varied the features of packet arrival processes are. Not only do users produce different types of workloads, but also, at a lower level, communication requests use a multitude of packet sizes and several different protocol types. In building models, one has to strike a balance between accuracy of feature representation on one side and simplicity and mathematical tractability on the other. However, often researchers opt for simplicity of mathematical analysis and ignore key features of their point processes. Although this approach can lead to useful approximations, unrealistic models are often obtained, which cannot be viewed and used with full confidence.

In this chapter we attempt to develop models of packet arrivals that are based on our knowledge of packet length and interarrival-time patterns present in our data and of certain statistical properties of these patterns. Packet length is not a factor in the Markov-Modulated Poisson-Process (MMPP) characterization of packet-arrival processes that was outlined in Chapter 4. The MMPP model is built on data for packet counts—groups of packets arriving in the same interval—and does not incorporate any data on packet length. However, packet length information is important if traffic models developed on the basis of our data are to be used in connection with studies of non-Ethernet networks, for variable packet lengths can have a major impact on traffic patterns in those networks. For instance, in an Asynchronous Transfer Mode (ATM) network, the transmission unit, the *cell*, has a constant length of 53 bytes, of which 48 are for the user's *payload*. Whereas a 46-byte Ethernet packet would correspond to a single ATM cell, a 1500-byte packet would require 32 ATM cells. Thus, for models of Ethernet workstation traffic to have more general applicability, they should include packet length information, and should model interarrival times rather than packet counts. Hence, we will seek two-dimensional point process models in which the coordinates of each point are an interarrival time and a packet length.

Our analysis of packet lengths in Chapter 2 indicated that packet lengths in our traces tend to be either shorter than 200 bytes or longer than 500 bytes. Furthermore, short packets were typically associated with long interarrival times (recall that we associate with an interarrival time the first packet of each packet pair generating the interarrival time), while long packets typically produced short interrival times. These patterns partition the packet-length and the interarrival-time spaces and suggest that an appropriate model might be one with two states: one for the short packets, and one for the long packets. (Notice that the two states in the MMPP model of Chapter 4 are really

only a mathematical construction and are not grounded on empirical reality.)

If the model is to have two states, we must determine how long each state should last. The answer to this question is provided by data on the lengths of sequences of short packets—strings of back-to-back short packets—and on the lengths of sequences of long packets. Section II examines our data on sequence length. We observe that the short- and long-packet sequence series are independent, and moreover, in each series, the probability density decreases approximately exponentially with the sequence length. These characteristics suggest the geometric distribution would model reasonably well the sequence lengths in each of the two states, and that a two-state Markov chain would be an appropriate model for the alternating process of short and long sequences.

Our next step is to propose a model that incorporates the Markov chain and produces interarrival times and packet lengths. We suggest that packet lengths and interarrival times can be generated in each state using probability density functions estimated from the data. Were these probability functions sampled independently, we would have a semi-Markov process. Since this is not the case with our processes, in Section III we extend the two-state semi-Markov process to the case in which the probability functions in the two states are not independent. We solve the model in terms of the autocovariances of the two component functions and fit it with the measured data. Unfortunately, the model's covariance does not approximate well the covariance of the data for intermediate and large values of the lag.

In order to understand the reasons for this divergence, we next generate several artificial point processes each of which contradicts one of the generalized model's assumptions. We discover that the problems are caused by the nonstationary components that are present in the measured arrival processes. Because of this negative result, we are forced to conclude that the extended semi-Markov model can only be used to represent accurately packet arrival processes over short time scales. However, we observe that the queuing dynamics of interest in interactive distributed systems, such as networks of personal workstations, occurs over short time scales.

II. Models for Packet Sequences

In Figure 3.15 we show the distribution of packet-lengths for aggregate traffic in the Sun Engineering network. We see that, with almost no exception, packets are either shorter than 200 bytes or longer than 500 bytes. Short packets are produced either by TCP character traffic, as discussed in Appendix A, or by Remote Procedure Call (RPC) packets, as discussed in Chapters 2 and 3. In contrast to the Berkeley network described in Appendix A, in which character traffic represented a sizable proportion of the total traffic (see Figure A.3), in the Sun Engineering network character traffic accounts for only about 6 percent of the total (see Figure 3.15). Instead, RPC packets account for more than 34 percent of the total packet traffic in the Sun network. The majority of long packets are 1500-byte fragments of larger messages generated by higher protocol layers. Fragmentation, involving a straightforward sequence of operations that can be streamlined, results in short interarrival times between fragments. As we have shown in Figures 4.10 and 4.11, when two successive messages are fragmented, the interarrival time between the last fragment of the first message and the first fragment of the second is normally large. (We have defined these pairs of fragments as ``LastFragment/1500''.)

Thus, since the distribution of packet lengths can be divided into two regions each of which has individual interarrival time properties and is produced by different system components, it is reasonable to classify packet lengths into two classes, *small* and *large*, depending on which of the two groups they belong to.

We will now compute the distribution of the lengths of short- and long-packet sequences—strings of consecutive short and strings of consecutive long packets. Following the methodology of Chapter 3, we will limit the data set only to pairs of packets with interarrival times of less than 80 milliseconds (as usual, each interarrival time will be associated with the first of the two packets that produced it). In Figure 5.1 we show the histogram of short-packet sequence lengths for the send queue of *station 25*; in Figure 5.2, the histogram of long-packet sequence lengths for the same station. Figures 5.3 and 5.4 display the same type of histograms for *station 25*'s receive queue. Both pairs of graphs are semi-logarithmic.

FIGURE 5.1. HISTOGRAM OF SHORT-PACKET SEQUENCE LENGTHS *STATION 25* – SEND QUEUE

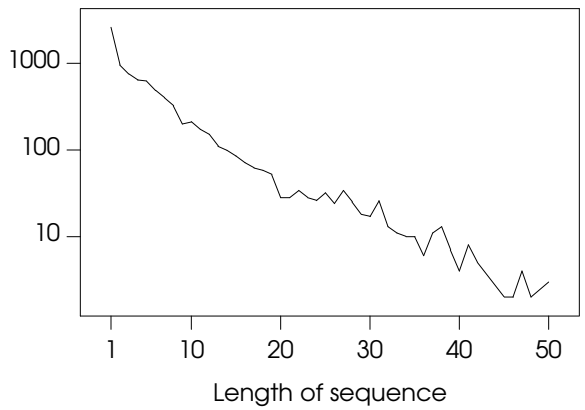

FIGURE 5.2. HISTOGRAM OF LONG-PACKET SEQUENCE LENGTHS *STATION 25* – SEND QUEUE

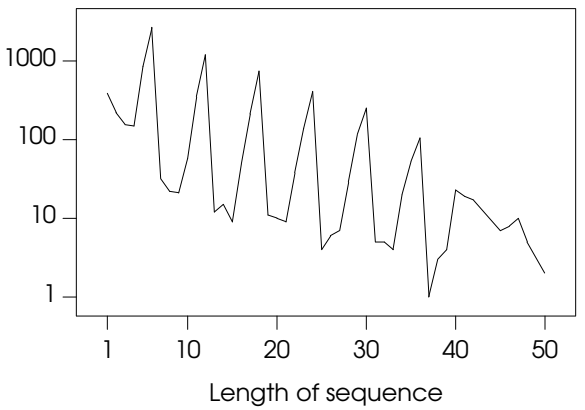

FIGURE 5.3. HISTOGRAM OF SHORT-PACKET SEQUENCE LENGTHS *STATION 25* – RECEIVE QUEUE

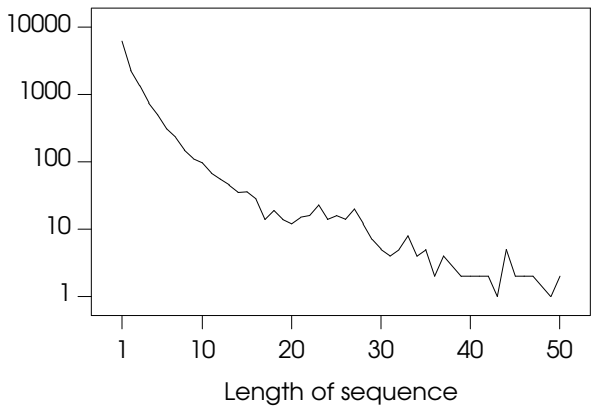

FIGURE 5.4. HISTOGRAM OF LONG-PACKET SEQUENCE LENGTHS *STATION 25* – RECEIVE QUEUE

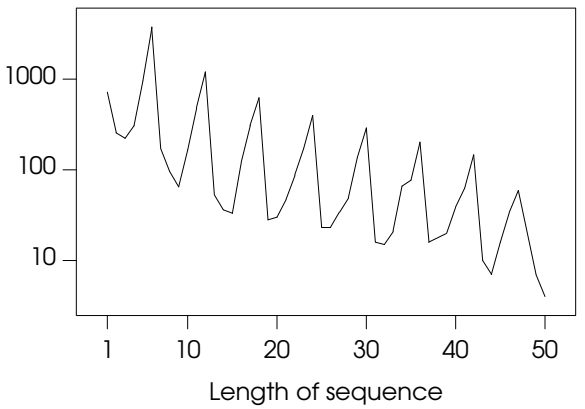

In Figures 5.2 and 5.4 we can clearly see the effects of the 8-Kbyte message fragmentation, which produces fragments in groups of six. Often, two or more groups follow each other producing peaks at values that are multiples of six. However, since we are

only considering interarrival times shorter than 80 milliseconds, in some cases we discard longer interarrival times produced by ``1500/LastFragment'' pairs (see Figures 4.10 and 4.11). When this happens, the lengths of some sequences of long packets are no longer a multiple of six. The existence of a few non-NFS sources of large packets and of concurrency among streams produced by independent processes running on the client workstation explains why Figures 5.2 and 5.4 do not consist simply of single lines at values that are a multiple of six.

The graphs in Figures 5.1 and 5.3 are approximated well by straight lines. Since these histograms are plotted as semi-log graphs, in which straight lines correspond to exponentially decreasing curves, the associated functions are negative exponentials. Although the graphs in Figures 5.2 and 5.4 show a more complex structure, the peaks nearly fall into a straight line and so do moving-average plots of the two graphs. Hence, we decide to ignore the level of detail represented by the peaks and to consider also the sequences of long packets as produced by exponentially decreasing probability density functions.

In Figures 5.5 and 5.6 we show the autocorrelation coefficients from lag 1 to lag 100 for the series of short-packet and long-packet sequence lengths, respectively, of *station 25*'s send queue. The dashed lines indicate the 99 percent confidence bands for uncorrelated series with the same number of elements. We see that the autocorrelation level is very small, and, for practical purposes, both series can be considered self uncorrelated.

FIGURE 5.5. CORRELOGRAM FOR THE SERIES OF LENGTHS OF SHORT PACKET SEQUENCES – *STATION 25*'S SEND QUEUE

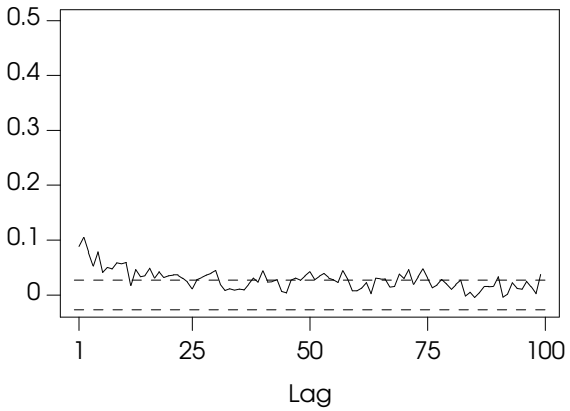

FIGURE 5.6. CORRELOGRAM FOR THE SERIES OF LENGTHS OF LONG PACKET SEQUENCES – *STATION 25*'S SEND QUEUE

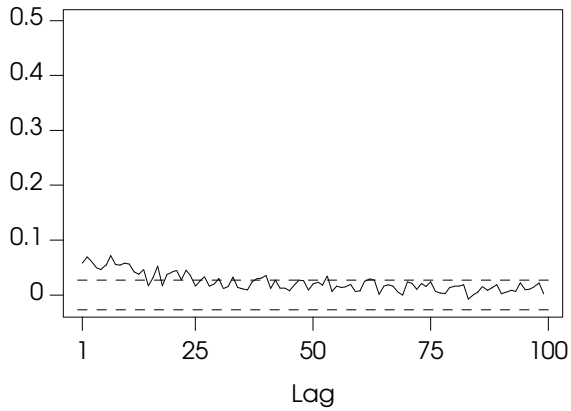

To confirm that there is no autocorrelation, we show in Figures 5.7 and 5.8 scatter plots at lag 1 of the lengths of short- and long-packet sequences, respectively. Each point in these graphs has for coordinates two successive sequence lengths, L_i on the x-axis and L_{i+1} on the y-axis. Although in these figures we have truncated the range of possible lengths at a maximum of 50 packets, in doing so we have dropped less than 0.5 percent of all sequences. Since sequences have discrete values, for clearer visualization of the point distribution we have added a random displacement to the coordinates of each point so that a point will lie anywhere in a square of size 1 centered on the original point position. The two plots are symmetric with respect to the line $y = x$, an indication

that the probability density functions of the sequence lengths have been sampled independently.

FIGURE 5.7. SCATTER PLOT OF LENGTHS OF SEQUENCES OF SHORT PACKETS – *STATION 25*'S SEND QUEUE

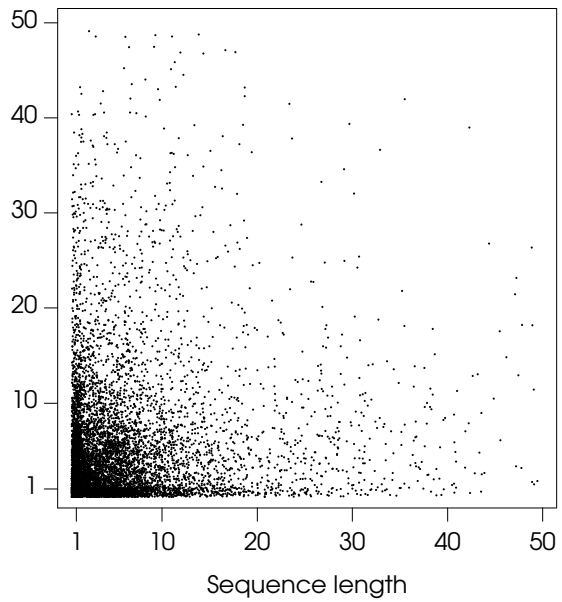

FIGURE 5.8. SCATTER PLOT OF LENGTHS OF SEQUENCES OF LONG PACKETS – *STATION 25*'S SEND QUEUE

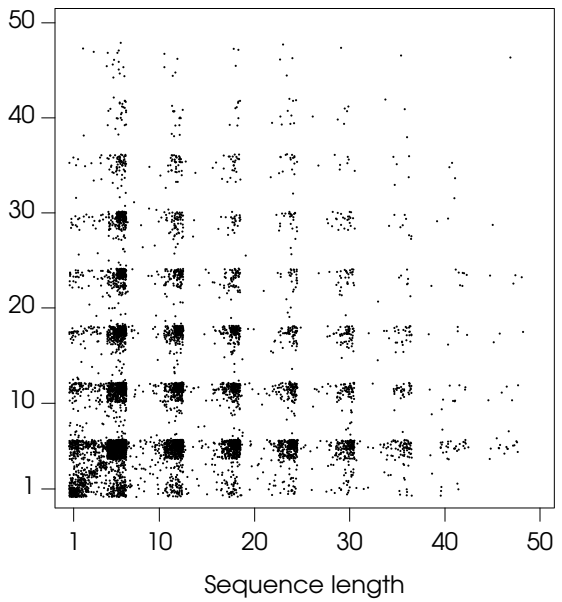

We conclude that the series whose histograms are shown in Figures 5.1 through 5.4 are produced by discrete probability functions that decrease roughly exponentially and are statistically independent within themselves. We have also verified that they are independent of each other. A geometric probability function should provide a good approximation for these data. However, since *station 25* is a diskless client workstation, before we can claim that geometric density functions are appropriate for modeling the short- and long-packet sequence lengths of our arrival processes, we must verify that client workstations with local disks and file servers also satisfy the assumptions of geometric probabilities.

In Figures 5.9 through 5.12 we show histograms of long- and short-packet sequence lengths for file server *station 21*. The histogram plots of the short-packet sequence lengths are characterized by steeply decreasing curves rather than straight lines in the region of small values. This higher proportion of short sequence lengths is caused by network contention. In the case of the receive queue, concurrent file accesses by different client workstations may result in a client inserting a large packet into a sequence of short packets transmitted by another client. In the case of the send queue, the concurrent activity is caused by the various server processes that handle client requests. The presence of multiple server processes produces asynchronous traffic, which again results in the intermixing of long and short packets addressed to different clients.

Asynchronous processes also play a role in the packet arrival processes of client workstations. Workstation users run multi-window interfaces, and in each window they run separate processes that access the file server independently and asynchronously

FIGURE 5.9. HISTOGRAM OF SHORT-PACKET SEQUENCE LENGTHS *STATION 21*'S – SEND QUEUE

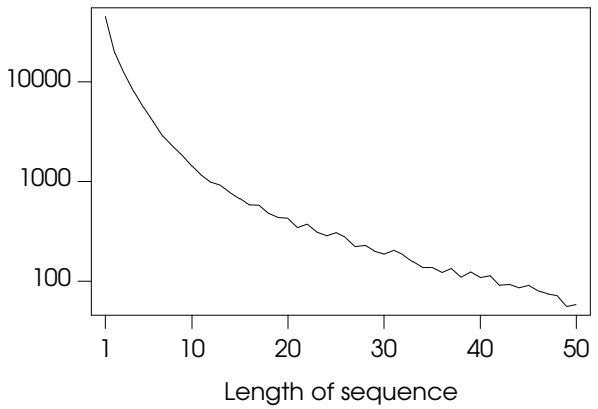

FIGURE 5.10. HISTOGRAM OF LONG-PACKET SEQUENCE LENGTHS *STATION 21*'S – SEND QUEUE

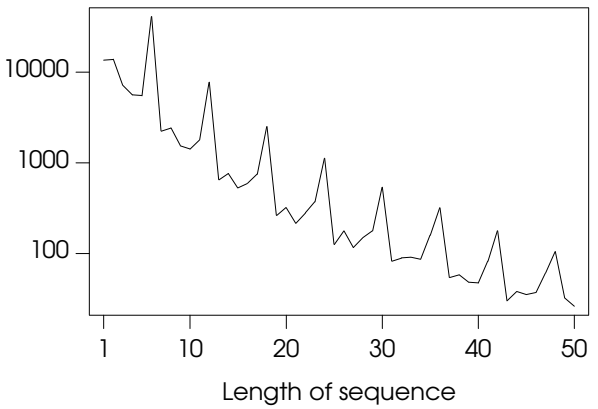

FIGURE 5.11. HISTOGRAM OF SHORT-PACKET SEQUENCE LENGTHS *STATION 21*'S – RECEIVE QUEUE

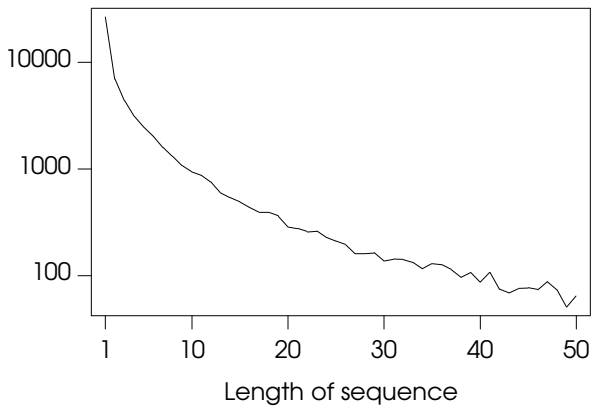

FIGURE 5.12. HISTOGRAM OF LONG-PACKET SEQUENCE LENGTHS *STATION 21*'S – RECEIVE QUEUE

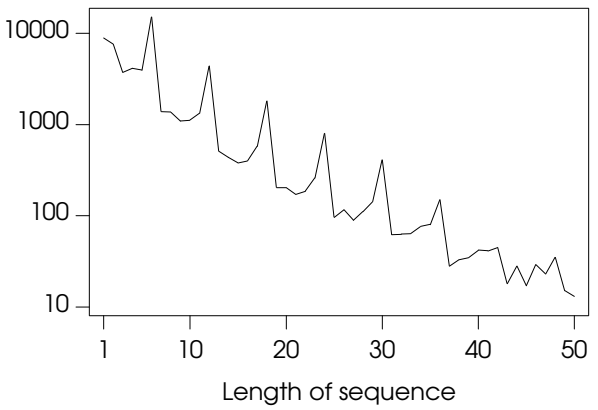

with respect to each other. Returning to Figures 5.1 and 5.3 for a moment, we see that in fact the curves in these two graphs also deviate slightly from a straight line in the short length region.

Figures 5.10 and 5.12 prompt two observations: first, these two curves (or, more precisely, their peaks) follow very nearly a straight line. Second, in comparison with Figures 5.2 and 5.4, there is a large proportion of sequence lengths occurring at values other than multiples of six. This is explained by the concurrency among different clients simultaneously communicating with the file server.

In Figures 5.13 through 5.16 we show the histograms of short- and long-packet sequence lengths of *station 11*, a client workstation with local disks. From these plots it is obvious that the ``straight line'' model does not apply to this machine. We have already observed in Chapter 3 (see Sections 3.IV and 3.VII) that *station 11* has a highly nonstationary behavior; unfortunately, such nonstationarity characterizes the interarrival times of most clients with local disks. Local disks are typically used as a swap partition and as a cache for frequently accessed files. Hence, if they are sufficiently large, these disks can reduce significantly the network traffic, resulting in sparse, highly irregular arrival

processes.

FIGURE 5.13. HISTOGRAM OF SHORT-PACKET SEQUENCE LENGTHS *STATION 11* – SEND QUEUE

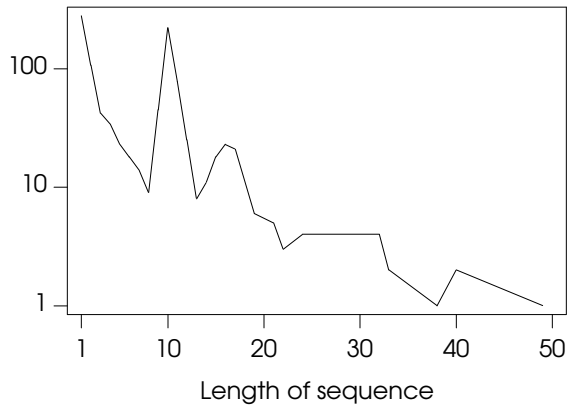

FIGURE 5.14. HISTOGRAM OF LONG-PACKET SEQUENCE LENGTHS *STATION 11* – SEND QUEUE

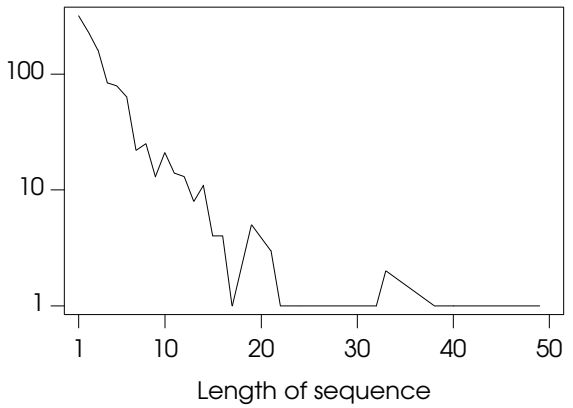

FIGURE 5.15. HISTOGRAM OF SHORT-PACKET SEQUENCE LENGTHS *STATION 11* – RECEIVE QUEUE

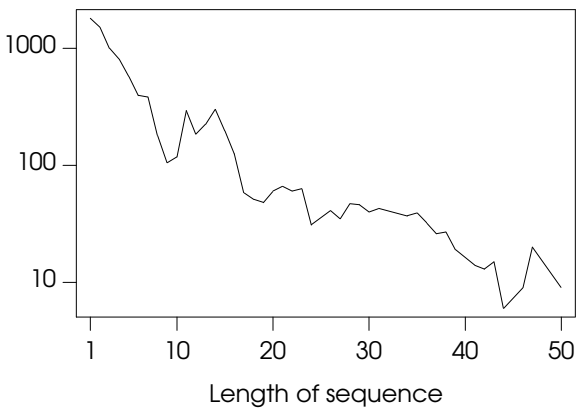

FIGURE 5.16. HISTOGRAM OF LONG-PACKET SEQUENCE LENGTHS *STATION 11* – RECEIVE QUEUE

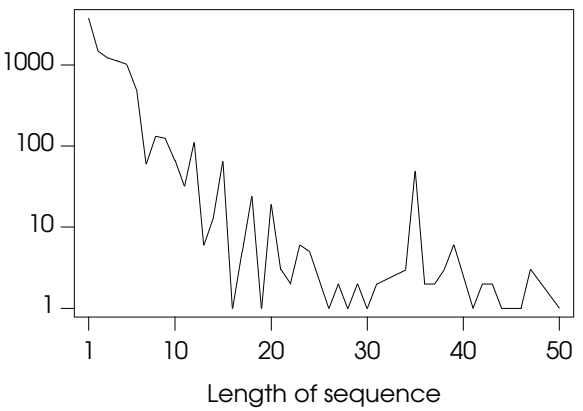

The fundamental difference between packet arrival processes generated by client workstations with local disks and those generated by diskless clients is that in the former the ratio of large packets to small packets is much lower than in the latter. Long packets, with their short interarrival times and low variability, tend to stabilize the arrival processes of diskless machines. The relatively large number of short packets, with their longer interarrival times and much higher variability, in the processes produced by disk-based clients, coupled with low packet arrival rates, results in unstable processes.

In Figures 5.17 and 5.18 we show the first 100 autocorrelation coefficients of short- and long-packet sequence lengths of the process generated by the *station 11*'s send queue. As in the other correlograms, we indicate within dashed lines the interval where with probability 0.95 the correlation coefficients of an uncorrelated process would lie. The confidence band is larger than in the previous cases because *station 11*'s send-queue arrival process is composed of many fewer arrivals. These graphs show that the sequence lengths of both short and long packets are essentially independent processes: the peak in Figure 5.17 around lag 63 is spurious. Indeed, it is produced by a small group

of relatively long sequences and by a single long sequence 63 sequences apart from the group. Since the sequences are independent and we have no other reason to believe that the straight line model would not (approximately) apply were *station 11*'s arrival process of higher intensity, we will consider this model valid also for clients with local disks.

FIGURE 5.17. CORRELOGRAM FOR THE SERIES OF LENGTHS OF SHORT PACKET SEQUENCES – *STATION 11*'S SEND QUEUE

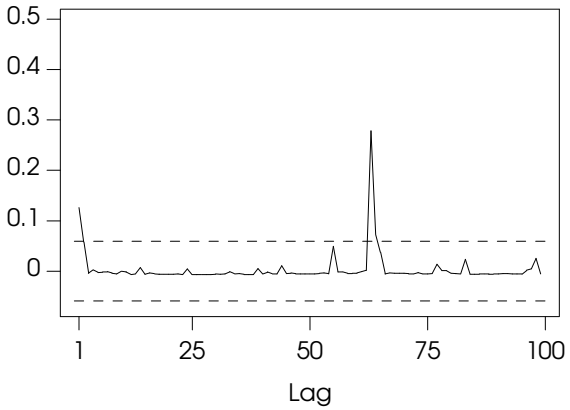

FIGURE 5.18. CORRELOGRAM FOR THE SERIES OF LENGTHS OF LONG PACKET SEQUENCES – *STATION 11*'S SEND QUEUE

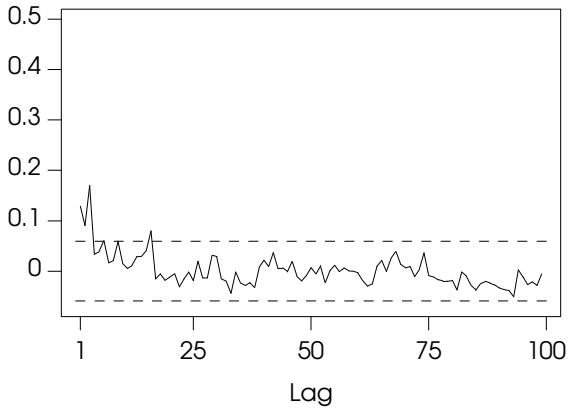

Given that sequences of short packets alternate with sequences of long packets, that the two types of sequences are independent of each other, and that both can be approximated by geometric distribution functions, the most natural choice for a mathematical model of these sequences is a two-state discrete-time Markov chain. We have estimated the transition probabilities for the processes generated by the send and receive queues of a number of workstations. These probabilities are shown in Tables 5.1 and 5.2.

TABLE 5.1. CONDITIONAL TRANSITION PROBABILITIES – SEND QUEUES

WORKSTATION	MATRIX ELEMENTS			
	p_{11}	p_{12}	p_{21}	p_{22}
station 2	0.865	0.135	0.076	0.924
station 3	0.868	0.132	0.079	0.921
station 7 (s)	0.910	0.090	0.075	0.925
station 8	0.887	0.113	0.063	0.937
station 11	0.993	0.007	0.242	0.758
station 13	0.910	0.090	0.082	0.918
station 14	0.828	0.172	0.091	0.909
station 15	0.913	0.087	0.048	0.952
station 16	0.946	0.054	0.072	0.928
station 20	0.901	0.099	0.184	0.816
station 21 (s)	0.840	0.160	0.136	0.864
station 22	0.927	0.073	0.019	0.981
station 23 (s)	0.753	0.247	0.071	0.929
station 25	0.838	0.162	0.085	0.915
station 26	0.966	0.034	0.020	0.980
station 27	0.984	0.016	0.004	0.996

We will use these transition probability matrices as the embedded two-state Markov chains of our semi-Markov models. One state corresponds to the generation of sequences of short packets, the other to the generation of sequences of long packets.

TABLE 5.2. CONDITIONAL TRANSITION PROBABILITIES – RECEIVE QUEUES

WORKSTATION	MATRIX ELEMENTS			
	p_{11}	p_{12}	p_{21}	p_{22}
station 2	0.789	0.211	0.065	0.935
station 3	0.793	0.207	0.058	0.942
station 7 (s)	0.968	0.032	0.077	0.923
station 8	0.954	0.046	0.038	0.962
station 11	0.936	0.064	0.256	0.744
station 13	0.863	0.137	0.073	0.927
station 14	0.674	0.326	0.063	0.937
station 15	0.938	0.062	0.028	0.972
station 16	0.673	0.327	0.019	0.981
station 20	0.971	0.029	0.117	0.883
station 21 (s)	0.916	0.084	0.115	0.885
station 22	0.883	0.117	0.041	0.959
station 23 (s)	0.848	0.152	0.095	0.905
station 25	0.691	0.309	0.063	0.937
station 26	0.890	0.110	0.027	0.973
station 27	0.978	0.022	0.014	0.986

In each state we will use two probability density functions to generate interarrival times and packet lengths. Our desire to define the models based on properties of the data has led us to modify the traditional semi-Markov model. This is discussed in the next section.

III. A Generalized Semi-Markov Process

The semi-Markov process was studied in the late fifties and early sixties by Smith [82] and Pyke [72,73], among others. A description of its main properties can be found in those studies. A semi-Markov process contains an embedded Markov chain that determines transitions between states. In each state intervals are chosen independently and identically distributed according to a probability density function associated with that state. The independence of successive intervals in each state makes the derivation of the properties of a semi-Markov process particularly simple. However, it also limits its suitability as a model for our arrival processes. Were we to use a two-state semi-Markov model, one state would represent the generation of short packets and the other the generation of long packets. However, as shown in Chapter 3, contrary to the assumptions on which the semi-Markov model is based, in each of these states interarrival times would be correlated.

Our solution is to use a modified two-state semi-Markov model, which we will term a *generalized semi-Markov* model, that does not assume independence of intervals in each state. The generalized semi-Markov model instead assumes that intervals in each state have a non-zero covariance. It consists of a two-state Markov chain with state-transition matrix

$$\mathbf{P} = \begin{bmatrix} \alpha_1 & 1-\alpha_1 \\ 1-\alpha_2 & \alpha_2 \end{bmatrix}, \tag{5.1}$$

and with equilibrium state probability vector $\boldsymbol{\pi} = (\pi_1, \pi_2)$

$$\pi_1 = \frac{1-\alpha_2}{2-\alpha_1-\alpha_2}, \qquad \pi_2 = \frac{1-\alpha_1}{2-\alpha_1-\alpha_2}. \tag{5.2}$$

In state 1, interarrival times are sampled from probability density function $f_1(x)$; in state 2, from probability density function $f_2(x)$. Each of the two series of intervals will be correlated; however, we continue to assume that intervals in one state are independent of intervals in the other state. In addition, we assume that the process generated in each state is only suspended, not disrupted, by a transition to the other state; upon a transition back to the original state, the process resumes at the point at which it was interrupted.

Under the assumptions stated above, the probability density function of interarrival times produced by the generalized semi-Markov process is

$$f_X(x) = \pi_1 f_1(x) + \pi_2 f_2(x) \ ,$$

and the mean interarrival time and the interarrival time variance are

$$E(X_i) = \pi_1\mu_1 + \pi_2\mu_2 \ ,$$

$$\text{var}(X_i) = \pi_1\sigma_1^2 + \pi_2\sigma_2^2 + \pi_1\pi_2(\mu_1 - \mu_2)^2 \ , \tag{5.3}$$

in which μ_1 and μ_2 are the means of the intervals generated in states 1 and 2, and σ_1^2 and σ_2^2 are the interval variances of the two states. We will next derive formulas for the covariance of intervals, correlation coefficients, and index of dispersion for intervals.

Let us use u to denote one state and $\bar{u}$ to denote the other state. If we consider two intervals X_i and X_{i+k} and denote with (v, w), where v and w are states, the fact that X_i is sampled from $f_v(x)$ and that X_{i+k} is sampled from $f_w(x)$, we have

$$E(X_iX_{i+k}\,|\,(u, u)) = \mu_u^2 + C_u^k \ ,$$

where C_u^k is a function, dependent on k, of the covariance of the intervals produced in state u. We also have

$$E(X_iX_{i+k}\,|\,(u, \bar{u})) = \mu_u\mu_{\bar{u}}$$

because intervals in one state are independent of the intervals in the other state.

We can express C_1^k and C_2^k in terms of the covariances of functions $f_1(x)$ and $f_2(x)$, which we will indicate with $\text{cov}_1(X_i, X_j)$ and $\text{cov}_2(X_i, X_j)$. Let us consider a sequence of $k+1$ Markov-chain states, $s_0s_1 \cdots s_k$, beginning and ending in the same state u, that is with $s_0 = u$ and $s_k = u$. Since we have assumed that $f_1(x)$ and $f_2(x)$ each evolve independently of any transitions to the other state, C_u^k will equal $\text{cov}_u(X_i, X_{i+m})$ with m equal to the number of occurrences of state u, in the Markov-chain sequence, minus 1. Hence, the range of m will be between 1 and k: 1 when the initial and the last states are the only occurrences of state u, and k when all the states in the sequence are u.

We will indicate with $P_{m,n}^{v,w}$ the probability that a sequence of length m, beginning in state v and ending in state w, contains n occurrences of state w. These probabilities can be easily expressed recursively:

$$P_{m,n}^{u,u} = \begin{cases} \alpha_u P_{m-1,n-1}^{u,u} & (n = m) \\ (1 - \alpha_u)P_{m-1,n-1}^{\bar{u},u} + \alpha_u P_{m-1,n-1}^{u,u} & (2 < n < m) \\ (1 - \alpha_u)P_{m-1,n-1}^{\bar{u},u} & (n = 2) \end{cases}$$

$$P^{\bar{u},u}_{m,n} = \begin{cases} (1-\alpha_{\bar{u}})P^{u,u}_{m-1,n} & (n = m-1) \\ \alpha_{\bar{u}}P^{\bar{u},u}_{m-1,n} + (1-\alpha_{\bar{u}})P^{u,u}_{m-1,n} & (1 < n < m-1) \\ \alpha_{\bar{u}}P^{\bar{u},u}_{m-1,n} & (n = 1) \end{cases}$$

with the following set of initial conditions

$$P^{1,1}_{2,2} = \alpha_1 \ , \quad P^{1,2}_{2,1} = 1-\alpha_1 \ , \quad P^{2,1}_{2,1} = 1-\alpha_2 \ , \quad P^{2,2}_{2,2} = \alpha_2 \ .$$

We can now express C^k_u in terms of the covariance of function $f_u(x)$

$$C^k_u = \sum_{q=1}^{k} P^{u,u}_{k+1,q+1}\,\mathrm{cov}_u(X_i, X_{i+q}) \ .$$

The unconditional probability that beginning in state v the Markov chain will be in state w after k steps is

$$\Pr(s_k = w,\ s_0 = v) = \Pr(s_k = w \,|\, s_0 = v)\Pr(s_0 = v) = p^k_{v,w}\pi_v \ .$$

We have indicated with $p^k_{v,w}$ the probabilities $\Pr(s_k = w \,|\, s_0 = v)$, that is, the elements of $\mathbf{P}^k$:

$$\mathbf{P}^k = \begin{bmatrix} \pi_1 + \pi_2\beta^k & \pi_2 - \pi_2\beta^k \\ \pi_1 - \pi_1\beta^k & \pi_2 + \pi_1\beta^k \end{bmatrix} ,$$

where $\beta = \alpha_1 + \alpha_2 - 1$. Then, the expectation $E(X_iX_{i+k})$ can be computed as

$$E(X_iX_{i+k}) = \pi_1 p^k_{1,1}\ (\mu_1^2 + C^k_1) + \pi_1 p^k_{1,2}\ \mu_1\mu_2 + \pi_2 p^k_{2,1}\ \mu_1\mu_2 + \pi_2 p^k_{2,2}\ (\mu_2^2 + C^k_2) \ ,$$

which can be expanded into

$$\begin{aligned} E(X_iX_{i+k}) = {} & \pi_1(\pi_1 + \pi_2\beta^k)\ \Big(\mu_1^2 + \sum_{q=1}^{k} P^{1,1}_{k+1,q+1}\,\mathrm{cov}_1(X_i, X_{i+q})\Big) + \\ & + \pi_1(\pi_2 - \pi_2\beta^k)\ \mu_1\mu_2 + \pi_2(\pi_1 - \pi_1\beta^k)\ \mu_1\mu_2 + \\ & + \pi_2(\pi_2 + \pi_1\beta^k)\ \Big(\mu_2^2 + \sum_{q=1}^{k} P^{2,2}_{k+1,q+1}\,\mathrm{cov}_2(X_i, X_{i+q})\Big) \ . \end{aligned}$$

We can now evaluate the covariance of the generalized semi-Markov process

$$\begin{aligned} \mathrm{cov}(X_i, X_{i+k}) = {} & E(X_iX_{i+k}) - E^2(X_i) = \\ & = (\mu_1 - \mu_2)^2\pi_1\pi_2\beta^k + \\ & \quad + \pi_1(\pi_1 + \pi_2\beta^k) \sum_{q=1}^{k} P^{1,1}_{k+1,q+1}\,\mathrm{cov}_1(X_i, X_{i+q}) + \\ & \quad + \pi_2(\pi_2 + \pi_1\beta^k) \sum_{q=1}^{k} P^{2,2}_{k+1,q+1}\,\mathrm{cov}_2(X_i, X_{i+q}) \ . \end{aligned} \tag{5.4}$$

Observe that the covariance depends not only on the covariances of the two component processes, but also on the difference between the means μ_1 and μ_2.

The series of correlation coefficients is obtained using the relationship $\rho_k = \mathrm{cov}(X_i, X_{i+k}) \,/\, \mathrm{var}(X)$, and the index of dispersion for intervals using Equation (4.3):

$$J_n = \frac{1}{\pi_1\mu_1 + \pi_2\mu_2}\Bigg\{\pi_1\sigma_1^2 + \pi_2\sigma_2^2 + \pi_1\pi_2(\mu_1 - \mu_2)^2 + \\ + 2\sum_{k=1}^{n-1}(1 - \frac{k}{n})\Bigg[(\mu_1 - \mu_2)^2\pi_1\pi_2\beta^k + \\ + \pi_1(\pi_1 + \pi_2\beta^k)\sum_{q=1}^{k} P_{k+1,q+1}^{1,1}\,\mathrm{cov}_1(X_i, X_{i+q}) + \\ + \pi_2(\pi_2 + \pi_1\beta^k)\sum_{q=1}^{k} P_{k+1,q+1}^{2,2}\,\mathrm{cov}_2(X_i, X_{i+q})\Bigg]\Bigg\}\,.$$

We will next attempt to verify the ability of the generalized semi-Markov model to capture the variability of our packet arrival processes.

IV. Data Fitting and Analysis

In Figure 5.19 we plot the first 100 autocorrelation coefficients of the interarrival-time series of *station 25*'s send queue and we also plot the autocorrelation coefficients produced by the generalized semi-Markov process using the covariances of *station 25*'s series. To compute the model's correlogram, we have separated *station 25*'s interarrival times into two series: one of interarrival times associated with short packets, the other of the interarrival times associated with long packets. (As mentioned earlier in Section I, we associate an interarrival time to the first of the two packets that produced it.) We have then computed the autocovariances of each series separately, and fed the coefficients into Equation (5.4). Although the two curves are close to each other for small values of the lag, there is a clear divergence between the model and the data points for lag values larger than 20: the data coefficients are above the 99 percent confidence band (indicated by the dotted line), whereas the model's coefficients are below it. Although the difference may seem negligible, in the calculation of the index of dispersion for intervals the effect of tiny coefficients is quite significant as a point at lag *j* on the IDI curve is constructed by adding together all of the correlation coefficients for lags smaller than *j*. To understand why data and model diverge, we will generate artificial component processes with controllable parameters, and compare the estimated autocorrelation coefficients of the process obtained combining the two component processes with the autocorrelation coefficients predicted by the generalized semi-Markov model.

We will construct the point processes from two component time series, which we assume are generated in state 1 and state 2 of the embedded Markov chain. We will produce autocorrelated interarrival times for the component time series using autoregressive functions of the form

$$Y_{i+1} = (1 - \gamma)g(\bar{\mu}) + \gamma Y_i\,, \tag{5.5}$$

where $g(\bar{\mu})$ is an exponential random variable with mean $\bar{\mu}$ and γ is between 0 and 1 [10]. In all of the examples in the remainder of this section we will compare the correlation coefficients of the simulated data with the correlation coefficients produced by the model. In order to focus on the effects of the correlations in the two component processes, we will generate components with the same mean so that the first term in Equation (5.4) will be zero.

FIGURE 5.19. AUTOCORRELATION COEFFICIENTS OF *STATION 25*'S SEND QUEUE AND MODEL

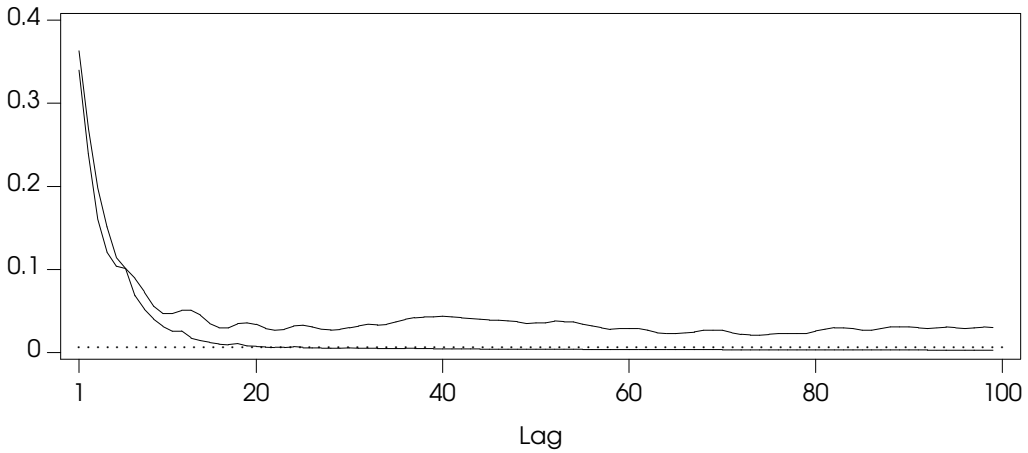

In Figure 5.20 we present a process that satisfies exactly all the model's requirements: we have two series that are mutually independent, each of the series is autocorrelated, the two series are mixed using a two-state Markov chain that determines the lengths of the sequences of interarrival times in each state, and transitions out of a state do not lose state memory. As we can see, there is a good match between the correlogram of the constructed process and that of the model evaluated using parameters from the component processes. (The small oscillations are caused by stochastic variations in the random number generator that produced the time series.)

To see whether correlations between the two components generate the difference seen in Figure 5.19, we evaluate the model using a process with two component time series that are mutually correlated. We generate one time series Y_i according to the autoregressive model shown above, and the second using the relationship

$$W_{i+1} = (1 - \theta) Y_i + \theta W_i \ .$$

We set the various coefficients so that the amount of cross correlations between the two component variables is significant. We show the first 100 correlation coefficients in Figure 5.22, where we also indicate the confidence band in which, with probability 0.95, the correlation coefficients would lie were the processes independent. In Figure 5.21 we graph the correlogram of the constructed process and that of the model evaluated using parameters from the component processes. Because the two curves are so close to each other, we conclude that the model is relatively independent of cross correlation between the two component functions.

We next test whether the clustering displayed in Figure 5.2 and 5.4 has an effect on the model's correlation structure. To do this, we generate a stream of interarrival times in which sequences of short packets are geometrically distributed and sequences of long packets are exactly six packets long. Within each sequence, the interarrival times are of the form (5.5). No major effect is discernible in the correlogram, shown in Figure 5.23; we, hence, conclude that packet clustering is not causing the divergence between the data and the generalized semi-Markov model in Figure 5.19.

FIGURE 5.20. DATA PRODUCED BY TWO FUNCTIONS SATISFYING ALL MODEL'S ASSUMPTIONS

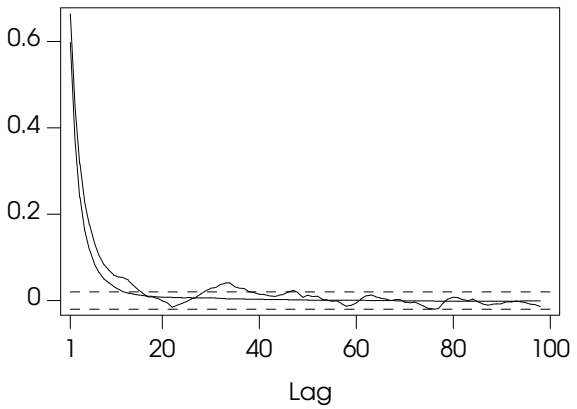

FIGURE 5.21. DATA PRODUCED BY TWO CROSS-CORRELATED FUNCTIONS

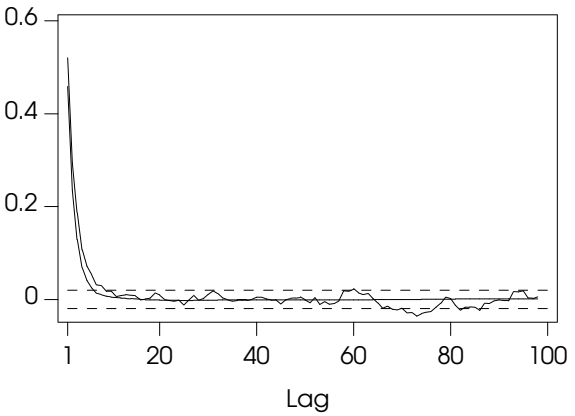

FIGURE 5.22. CROSS CORRELATION COEFFICIENTS OF COMPONENT PROCESSES

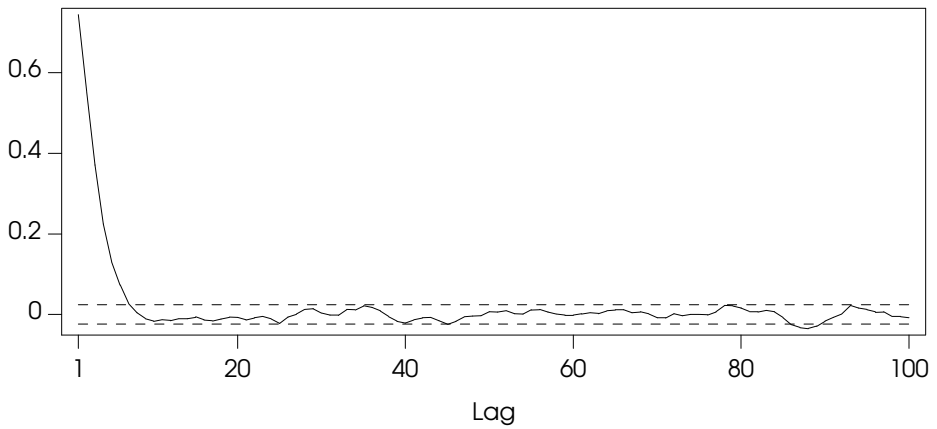

We, finally, experiment with nonstationary data in Figure 5.24. We produce two independent, moderately nonstationary series by slowly increasing the mean $\bar{\mu}$ in the autoregressive model (5.5). Now we see that the correlogram of the data decreases very slowly, and, indeed, lies outside the 95-percent confidence interval. As one could have guessed, the model's output is instead more regular and quickly positions itself inside the confidence band. (Recall that the confidence interval delimits the region where, in this case with probability 0.95, all the coefficients would be were not the process significantly autocorrelated.) As in Figure 5.19, here too what appears a small difference between the correlation coefficients is greatly amplified during the computation of the index of dispersion for intervals. Although we have produced these data with simple, linearly increasing trends, while, as discussed in Chapter 3, nonstationarities in our arrival processes cannot be classified as trends, the effects of packet arrival nonstationarities on the generalized semi-Markov model follow the same patterns.

FIGURE 5.23. DATA PRODUCED WHEN THE PROBABILITY DISTRIBUTION IN ONE OF THE STATES IS DETERMINISTIC

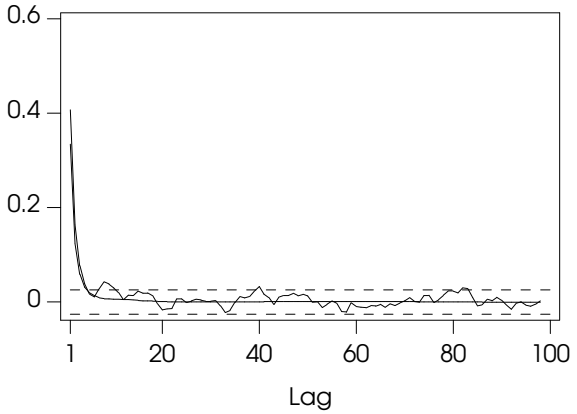

FIGURE 5.24. DATA PRODUCED BY NONSTATIONARY COMPONENT PROCESSES

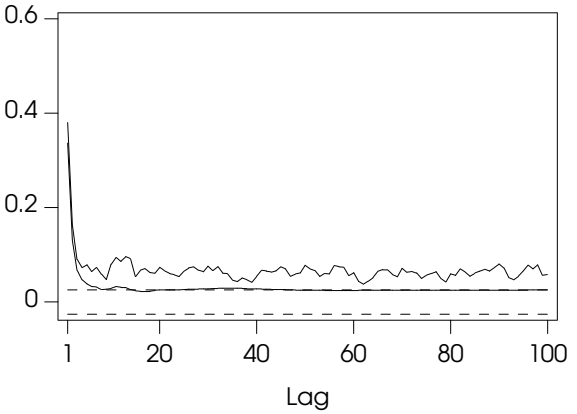

We have come to realize that much of the medium- and long-term variability of packet arrival processes, as revealed by the index of dispersion for intervals, is a direct effect of the processes' slowly decreasing autocorrelation coefficients, and is caused by nonstationary process components. In Section 3.V, we argued that is it rather difficult to isolate and remove these components from our arrival processes. As a result, we are forced to reject the generalized semi-Markov model as a model appropriate for the representation of packet arrival processes that are nonstationary over long time scales. Nevertheless, in the next section we will provide approximations, based on our extended semi-Markov process, that are valid for short-term variability.

V. Short-Interval Approximations

It is unfortunate, though not quite unexpected, that our generalization of the semi-Markov process performs poorly with nonstationary processes. However, in interactive systems such as the queuing systems found in workstation networks, the system behavior over short intervals is perhaps the key performance indicator. These systems are designed to be highly responsive, and queue lengths must be kept short. Hence, the queuing behavior is associated to the statistical interactions among packet arrivals for small lag values. Although the high values of the asymptotes in the index of dispersion curves shown in Chapter 4 indicate that interval and count processes are affected by considerable long-term variability, one could argue that long-term arrival-process fluctuations have no effects on queues if the system manages to keep queue lengths short. In addition, as shown in Figure 5.19, the model fit over small lags is acceptable (and more evident if we smooth the data autocorrelation coefficients). Therefore, we propose to use the generalized semi-Markov model to represent the variability of the arrival processes over short intervals. (Notice that a simpler semi-Markov process would not suffice as a suitable model since its autocorrelation coefficients are of the form Kb^n, with both K and b less than 1, which hence decrease too fast [18].) In order to fit the model to our arrival processes, in addition to the transition probabilities α_1 and α_2 of matrix (5.1) (that is p_{11} and p_{22} in Tables 5.1 and 5.2), we need to provide the mean and variance of the two component processes, as well as their autocovariances.

It is infeasible to express the correlation structure of the two component processes in terms of their autocovariance coefficients, as we have done in the development of the theory in Section III. For ease of use, here we will instead provide algebraic approximations based on correlation coefficients. (We chose the correlation coefficients for it is easier to work with their normalized scale; for each process, the autocovariances can be obtained by multiplying the autocorrelation coefficients by the variance of the process found in Table 5.3 below.) It turns out that the autocorrelation functions of both component processes are well approximated by straight lines, as the examples in Figures 5.25 and 5.26 show. The approximation holds for all types of workstations: diskless, with local disks, file servers. Figure 5.25 shows the autocorrelation coefficients of interarrival times between short packets and, as an almost horizontal line near 0, the autocorrelation coefficients of interarrival times between long packets for the send queue of client *station 3*. Figure 5.26 presents the same data for the two component processes produced by client *station 27*, a workstation with a local disk.

FIGURE 5.25. AUTOCORRELATION COEFFICIENTS OF PROCESS COMPONENTS OF *STATION 3*'S SEND QUEUE

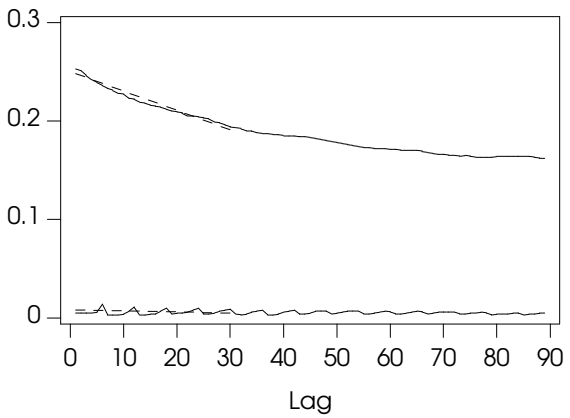

FIGURE 5.26. AUTOCORRELATION COEFFICIENTS OF PROCESS COMPONENTS OF *STATION 27*'S RECEIVE QUEUE

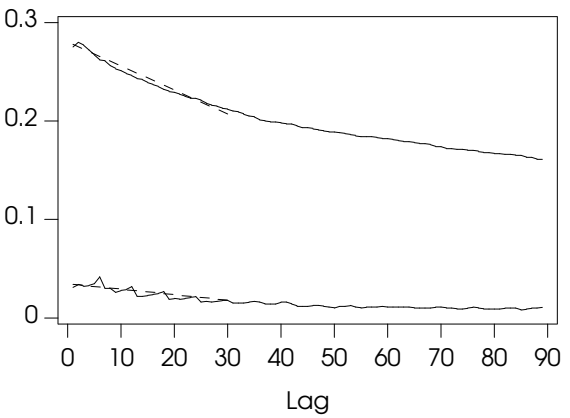

The linear approximations that we will tabulate for the autocorrelation coefficients of interarrival times in Tables 5.1 and 5.2 are valid only between lags from 1 to 30. (Notice that we have excluded the lag-0 coefficient, which is always 1.) The time spans of the linear models range from about 218 ms to 517 ms depending on the mean interarrival times of the various workstations. These intervals, outside which we suggest not to use our approximation models, may seem small. However, only very rarely do we observe queue lengths of size 30 or more. Hence, the dynamics of the queues of workstation systems should be captured by the generalized semi-Markov model truncated to these time scales.

Table 5.3 lists estimated values of the means and variances of the processes produced by the send queues of a representative set of workstations. As expected, in general faster architectures produce shorter mean interarrival times. The table also shows the estimated values of the mean and variance for each of the two component processes: the short-packet interarrival-time process, indicated by subscript 1, and the long-packet one, with subscript 2. Observe how much smaller the mean interarrival times of long packets are than those of short packets. However, their variances display a

much wider range than the variances of short-packet interarrival times. This is related to the long interarrival times that we have classified as ``LastFragment/1500''. As we can see, the variance is higher in machines with local disks, an indication that these workstations transfer longer data sets, and thus produce more ``LastFragment/1500'' interarrival times.

TABLE 5.3. MEANS AND VARIANCES OF MEASURED PROCESSES – SEND QUEUES
Means in ms Variances in ms^2 Errors in percentage

WORKSTATION	AGGREGATE PROCESS				COMPONENT PROCESSES			
	mean	mean error	var	var error	$mean_1$	var_1	$mean_2$	var_2
station 2	12.391	–0.044	325.983	–0.039	26.558	415.926	4.407	98.417
station 3	15.909	–0.028	414.982	–0.018	34.028	431.696	5.058	90.632
station 7	11.004	–5.878	231.408	–3.693	16.853	270.671	4.944	118.554
station 8	13.053	–0.297	347.386	–0.241	26.024	459.565	5.761	136.567
station 11	16.168	–0.373	323.939	0.087	16.348	322.958	7.799	296.829
station 13	14.683	–6.790	334.699	–4.438	24.069	377.069	4.226	79.991
station 14	12.598	0.237	339.101	0.209	28.599	457.337	4.178	71.276
station 15	12.406	0.281	408.470	0.038	22.112	453.597	7.105	303.740
station 16	17.251	–0.691	375.375	–0.302	27.420	368.034	3.414	53.207
station 20	14.312	–0.043	362.179	–0.045	19.039	396.539	5.509	178.828
station 21	11.816	9.165	244.905	5.694	20.366	291.204	6.552	143.674
station 22	9.005	0.131	313.545	0.161	31.185	729.676	3.247	44.675
station 23	8.552	15.094	229.607	12.800	25.264	352.720	5.410	144.048
station 25	12.559	0.778	346.452	0.678	29.262	468.933	3.944	65.182
station 26	10.152	–0.775	254.886	–0.669	20.852	372.369	3.733	74.529
station 27	7.278	–4.771	140.489	–3.363	14.778	183.706	4.969	104.536

Using the transition probabilities of Table 5.1 we can compute the equilibrium state probabilities (5.2), and with them the model's mean and variance from Equations (5.3). In the columns ``mean error'' and ``var error'' in Table 5.3, we list the percent errors on the model's predictions. Except for one workstation, the percent errors are in absolute value smaller than 10, proving again that the Markov property is satisfied in our data.

In order to evaluate the generalized semi-Markov models, we also need the autocovariances of the component processes. As mentioned earlier, we will provide first-order polynomial formulas (i.e. equations of straight lines) that interpolate accurately the first 30 autocorrelation coefficients. In order to obtain the covariances, one will have to multiply the correlation equations by the values of the (aggregate) process variances found in Table 5.3. In Table 5.4, we list the coefficients of interpolating linear equations of the form $y = mx + q$. (Again, with subscript 1 we label quantities that refer to short-packet processes.) If both coefficients are very small, a situation that indicates statistical independence between subsequent arrivals, we enter zeros in the table.

VI. Models of Packet Lengths

In order to complete our data characterization, which, as mentioned in the Introduction, is based on a two-dimensional model, we now describe model approximations for lengths of small and large packets. Several observations are in order. First, large NFS packets have a very simple structure: they come in groups of six, the first five fragments in a group are 1500 bytes long and the last fragment is somewhat shorter, for instance 920 bytes. Groups can follow each other, but as shown in Figures 5.2 and 5.4 and as discussed in Section II, the probability that k groups occur consecutively decreases

TABLE 5.4. CORRELATION COEFFICIENT INTERPOLATION LINES – SEND QUEUES

WORKSTATION	EQUATION COEFFICIENTS			
	m_1	q_1	m_2	q_2
station 2	-0.00218	0.193	0	0
station 3	-0.00204	0.255	0	0
station 7	-0.00229	0.19	-0.000179	0.0154
station 8	-0.00368	0.159	-0.000179	0.0214
station 11	-0.00025	0.235	-0.000107	0.0242
station 13	-0.00279	0.227	0	0
station 14	-0.00211	0.141	-0.000107	0.0202
station 15	-0.00164	0.342	-0.00196	0.187
station 16	-0.00114	0.121	-0.000393	0.0188
station 20	-0.000357	0.0907	-0.00132	0.144
station 21	-0.00261	0.169	0	0
station 22	-0.00129	0.548	-0.000214	0.0214
station 23	-0.00129	0.0936	0	0
station 25	-0.00225	0.117	-0.000536	0.0221
station 26	-0.00179	0.175	-0.000429	0.0459
station 27	-0.00457	0.328	0	0

geometrically as *k* increases. This structure produces correlation at values that are multiples of six.

Second, the other components in the distribution of large packets, TCP and ND packets, which are shown in Figure 2.15, at 1064 and 1072 bytes respectively, are irrelevant. ND, the Network Disk protocol, is being phased out and, although it produces about 4 percent of the packet count (including also small packets), among the workstations that we have considered only *station 1* uses it. Long TCP packets account for about 3 percent of the total, but they also occur in groups and hence, rarely mix with NFS packets. We will ignore both TCP and ND packets.

Third, some complications occur with file servers as their traffic is the superposition of traffic simultaneously directed to or coming from several client workstations. This will have the effect of making the lengths of sequences of 1500-byte packets variable, as discussed in Section II with reference to Figures 5.10 and 5.12. As a result of these considerations, we propose three possible models, with varying degrees of approximation, for long packet sizes: the first model would produce one 920-byte packet every five 1500-byte packets; the second model is a mean/standard-variation model with mean 1373 bytes and standard variation 247 bytes; and the third model, a special case of the second, would produce a constant packet of size 1373.

Fourth, the structure of the short-packet distribution is rather complicated. The sequence of short-packet lengths is strongly autocorrelated, as contiguous packets tend to have similar lengths. However, these correlations lose importance if they are put in perspective. In Figure 5.27 we reproduce the aggregate packet length distribution on a linear scale. From this graph we can see that the autocorrelation within the series of short packets, which are all concentrated at one end of the figure, amounts to small variations in packet sizes when viewed from the other end. Thus, we will ignore the fine structure produced by occurrences of groups of short packets of similar lengths, and describe short packets simply by their means and variances. In most cases there are no significant differences among the various workstations, hence we will simply provide the mean short packet and the short-packet standard variation based on the aggregate traffic. They

FIGURE 5.27. PERCENTAGE OF PACKETS VS. PACKET LENGTH

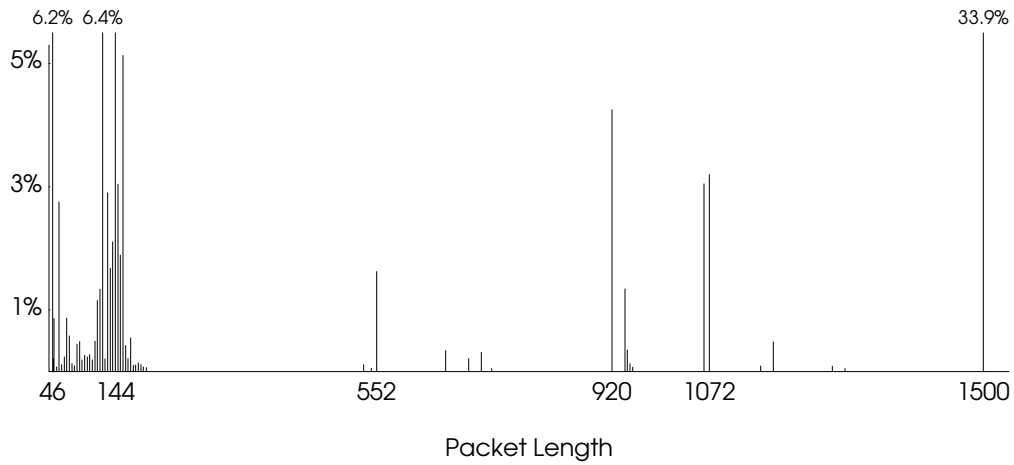

are 119 bytes and 50 bytes, respectively.

VII. Summary

In this chapter we have constructed models of packet arrival processes using properties of the data. We focused on interarrival times, since we felt that packet lengths needed to be represented explicitly, which is not possible packet count models. We found that sequences of short and long packets are modeled accurately by two-state discrete-time Markov chains. In order to represent correlations among arrivals in each of the states, we have extended the semi-Markov model and derived formulas for the autocovariances, and index of dispersion of intervals.

Unfortunately, the generalized semi-Markov model does not reproduce accurately the autocorrelation coefficients for medium and large lag values. Nonetheless, we provided model approximations based on linear interpolation of autocorrelation coefficients of the component processes that are valid and accurate for small lag values.