A GSMP Formalism for Discrete Event Systems

PETER W. GLYNN

Invited Paper

We describe here a precise mathematical framework for the study of discrete event systems. The idea is to define a particular type of stochastic process, called a generalized semi-Markov process (GSMP), which captures the essential dynamical structure of a discrete event system. The paper also attempts to give a flavor of the qualitative theory and numerical algorithms that can be obtained as a result of viewing discrete event systems as GSMPs.

I. Introduction

A fundamental obstacle to the study of discrete event systems is the lack of a comprehensive framework for the description and analysis of such systems. In this paper, we attempt to give such a framework.

The idea that we shall pursue here is to define a particular type of stochastic process, called a generalized semi-Markov process (GSMP), which captures the essential dynamical structure of a discrete event system. We view the GSMP framework both as a precise "language" for describing discrete event systems, and as a mathematical setting within which to analyze discrete event processes.

We start, in Section II, by giving an abstract description of a discrete event system. At this level of abstraction, some of the connections between continuous variable dynamic systems (CVDSs) and discrete event dynamic systems (DEDSs) become clear; this discussion is in the spirit of [9]. Section III specializes the above abstract framework by specifying a GSMP as a particular type of event-driven stochastic process. The GSMP structure is then immediately applied to develop a variance reduction technique that is potentially applicable to a vast array of discrete event simulations.

In Section IV, the GSMP framework is specialized still further, thereby yielding the class of time-homogeneous GSMPs. These processes can be analyzed via Markov chain techniques. These Markov chain ideas are then exploited in order to obtain some qualitative results pertaining to the "long-run" behavior of discrete event systems. Section V explores the relationship between continuous-time Markov chains, semi-Markov processes, and GSMPs.

Section VI returns to time-homogeneous GSMPs, this time exploring the qualitative theory from a regenerative process point of view. A new type of regenerative structure is described for discrete-event systems which are scheduled by "exponentially bounded" distributions. This technique can be applied to more general GSMPs with some additional work, but we present here only the current version. With the aid of regenerative process ideas, a strong law and a central limit theorm for discrete event systems are established. In our opinion, these results are typical of what we can expect to hold for "well-behaved" discrete event processes.

Finally, in Section VII, we give a flavor of the computational enhancements to discrete event simulations that are possible by making explicit use of the GSMP framework. Specifically, likelihood ratio ideas for importance sampling are briefly described.

Manuscript received August 24, 1988; revised September 15, 1988. This research was supported under U.S. Army Research Contract DAAG-84-K-0030, and National Science Foundation Grant DCR-8509668.

The author is with the Department of Operations Research, Stanford University, Stanford, CA 94305-4022, USA.

IEEE Log Number 8825173.

II. The CVDS/DEDS Analogy

Suppose that we wish to model the output process $(s(t): t \geq 0)$ corresponding to a (deterministic) CVDS. Frequently, the approach taken is to try to represent $s(t)$ in the form $s(t) = h(x(t))$, where $x(t)$ is some suitably chosen characterization of the "internal state" of the system. Thus, given the output process $(s(t): 0 \leq s \leq T)$, we can extend the output process to the interval $(T, T + h]$ by computing the internal state x over the interval and setting $s(t) = h(x(t))$, $T < t \leq T + h$.

The typical approach used to model a DEDS is similar in concept. We first recall that DEDSs are frequently used as models of systems having piecewise constant trajectories. For example, the trajectory of a queueing system is constant between arrival and departure epochs of customers. As a consequence, if $(S(t): t \geq 0)$ is the output process corresponding to a discrete event system, it typically takes the form

$$S(t) = \sum_{n=0}^{\infty} S_n I(\Lambda(n) \leq t < \Lambda(n+1)) \qquad (2.1)$$

0018-9219/89/0100-0014$01.00 © 1989 IEEE

where we require that $0 = \Lambda(0) < \Lambda(1) < \cdots$. ($I(A)$ represents an indicator function which is one or zero depending on whether or not A occurs.) In the representation (2.1), S_n represents the output state at the nth transition epoch, and $\Lambda(n)$ is the instant at which the nth transition occurs. Then, $\Delta_{n+1} = \Lambda(n + 1) - \Lambda(n)$ is the time between the nth and $(n + 1)$th transitions of $(S(t): t \geq 0)$.

To characterize the dynamics of the output process $(S(t): t \geq 0)$, we assume the existence of a stochastic sequence $X = (X_n: n \geq 0)$ which describes the time-evolution of the internal state of the system. (We permit X to be stochastic in order to allow the discrete event system to have random behavior.) We require that the (S_n, Δ_n)'s be related to the internal state sequence X via a mapping of the form $(S_n, \Delta_n) = (h_1(X_n), h_2(X_n))$. Given the output process $(S(t): 0 \leq t \leq \Lambda(n))$, we can then extend the output process to the interval $(\Lambda(n), \Lambda(n + 1)]$ by first computing X_{n+1}. We then calculate $\Delta_{n+1} = h_2(X_{n+1})$, $\Lambda(n + 1) = \Lambda(n) + \Delta_{n+1}$, and set $S(t) = S(\Lambda(n))$ for $\Lambda(n) < t < \Lambda(n + 1)$. We complete the extension to $(\Lambda(n), \Lambda(n + 1)]$ by calculating $S_{n+1} = h_1(X_{n+1})$ and setting $S(\Lambda(n + 1)) = S_{n+1}$. This recursive approach to defining $(S(t): t \geq 0)$ works, provided that the output process is *nonexplosive* (i.e., $\Lambda(n) \to \infty$ a.s. as $n \to \infty$).

The above discussion shows that both the CVDS and DEDS approaches to modeling the output of a system are, in principle, solved, once we characterize the internal state of the system. For a CVDS, perhaps the most general characterization is to assume that the internal state process $x = (x(t): t \geq 0)$ satisfies a relation of the form

$$x = f(x) \tag{2.2}$$

for some mapping f. In other words, for each $t \geq 0$, this formulation requires specifying a mapping $f(t, \cdot)$ for which

$$x(t) = f(t, x(s): 0 \leq s < \infty) \tag{2.3}$$

must hold. The analogous condition for a DEDS is to require that there exist a sequence of independent r.v.'s $\eta = (\eta_n: n \geq 0)$, and a map f such that

$$X = f(X, \eta). \tag{2.4}$$

This is equivalent to requiring the existence of a family of (component) mappings $f_{n+1}(\cdot)$ such that

$$X_{n+1} = f_{n+1}(X_k, \eta_k: 0 \leq k < \infty). \tag{2.5}$$

The resemblance between (2.2)–(2.3) and (2.4)–(2.5) should be clear.

Of course, the noncausal nature of (2.3) creates difficulties both mathematically and computationally. Furthermore, formulation of a model for x directly in terms of the mappings $(f(t, \cdot): t \geq 0)$ is often unnatural. As a consequence, it is more typical to limit models of the internal state process $x = (x(t): t \geq 0)$ to relations of the form

$$\text{s.t.} \quad \begin{aligned} x'(t) &= f(t, x(s): 0 \leq s \leq t) \\ x(0) &= x_0, \end{aligned} \tag{2.6}$$

for some prescribed family of mappings $(f(t, \cdot): t \geq 0)$, and initial condition x_0. (Note that we are now implicitly assuming that the internal state takes values in $\mathbb{R}^d$.) The model (2.6) is causally defined in terms of the infinitesimal characteristics of the system. The "local specification" that is implicit in asserting that x satisfies a given differential equation is generally easier to formulate from a modeling point of view than the "global specification" implicit in (2.3). Note that the representation (2.6) permits x to be described by differential equations with delay, as well as certain types of integro-differential equations.

Of course, the analog to the causal representation (2.6) for a DEDS is to assume that the internal state sequence X satisfies

$$\text{s.t.} \quad \begin{aligned} X_{n+1} &= f_{n+1}(\eta_{n+1}, X_k: 0 \leq k \leq n) \\ X_0 &= x_0. \end{aligned} \tag{2.7}$$

In Section III, we describe a family of discrete event systems in which X has the general form given by (2.7).

A limitation of the causal models (2.6) and (2.7) is that the mathematical theory available to study such unstructured systems is rather poorly developed. Fortunately, in many applications of CVDS, it is possible to choose the internal state process x so that the dynamics are described by a differential equation of the form

$$\text{s.t.} \quad \begin{aligned} x'(t) &= f(t, x(t)) \\ x(0) &= x_0. \end{aligned} \tag{2.8}$$

As we shall indicate shortly, the mathematical theory pertinent to (2.8) is quite extensive.

The DEDS analog of the representation (2.8) is to require that the internal state sequence X satisfy a recursion of the form

$$\text{s.t.} \quad \begin{aligned} X_{n+1} &= f_{n+1}(X_n, \eta_{n+1}) \\ X_0 &= x_0. \end{aligned} \tag{2.9}$$

Such representations can often be obtained for systems satisfying (2.7), provided that a judicious choice of state space is made. Specifically, it is often possible to adjoin "supplementary variables" χ_n to a state descriptor X_n satisfying (2.7), to obtain a new state sequence $X'_n = (X_n, \chi_n)$ satisfying (2.9).

The mathematical power of the representation (2.9) is a consequence of the following easily proved result.

Proposition 1: Suppose that X satisfies (2.9), and that the r.v.'s $\{x_0, \eta_n: n \geq 1\}$ are independent. Then, X is a Markov chain (i.e., $P\{X_{n+1} \in \cdot \mid X_0, \cdots, X_n\} = P\{X_{n+1} \in \cdot \mid X_n\}$).

A substantial literature on the theory of Markov chains can be applied to the analysis of DEDS for which the internal state sequence satisfies the conditions of Proposition 1. Similarly, the vast mathematical theory on differential equations is directly relevant to CVDS satisfying the representation (2.8). For example, existence theory for the differential equation (2.8) basically yields conditions under which there exists an output process (compatible with (2.8)) which can be defined over the entire semi-infinite interval $[0, \infty)$. The DEDS counterpart involves deriving conditions under which $(S(t): t \geq 0)$ is nonexplosive.

Much of the differential equations literature on systems satisfying (2.8) pertains to systems obeying the stronger condition

$$\text{s.t.} \quad \begin{aligned} x'(t) &= f(x(t)) \\ x(0) &= x_0. \end{aligned} \tag{2.10}$$

This literature typically focuses on the large-time behavior of the internal state process $(x(t): t \geq 0)$. This, in turn, is strongly related to the study of the set $\{x: f(x) = 0\}$ of equilibrium points for (2.10).

The DEDS counterpart to (2.10) requires a model formulation in which the internal state sequence X takes the form

$$\begin{aligned} X_{n+1} &= f(X_n, \eta_{n+1}) \\ \text{s.t.} \quad X_0 &= x_0. \end{aligned} \tag{2.11}$$

The following result is easily demonstrated, and so the proof is omitted.

Proposition 2: Suppose that X satisfies (2.11). In addition, assume that $\{\eta_n: n \geq 1\}$ is a collection of independent identically distributed (i.i.d.) r.v.'s, independent of x_0. Then, X is a time-homogeneous Markov chain (i.e., there exists a transition function $P(x, A)$ such that $P\{X_{n+1} \in \cdot \mid X_0, \cdots, X_n\} = P(X_n, \cdot)$).

As in the CVDS setting, much of the mathematical literature on Markov chains of the form (2.11) concentrates on study of the long-run behavior of the system. The concept of equilibrium point is now replaced by that of an invariant probability distribution. A probability distribution π is said to be *invariant* for the (time-homogeneous) Markov chain X if

$$\pi(dy) = \int_{\Sigma} \pi(dx)\, P(x, dy) \tag{2.12}$$

(Σ is the state space of X). In the presence of irreducibility hypotheses on X, the existence of an invariant probability distribution π typically implies that for each (measurable) subset A of Σ

$$\frac{1}{n} \sum_{k=0}^{n-1} I(X_k \in A) \rightarrow \pi(A) \quad \text{a.s.}$$

as $n \rightarrow \infty$, for every possible initial condition x_0. The analogous CVDS concept is that of (global) stability: there exists $\bar{x}$ such that $x(t) \rightarrow \bar{x}$ as $t \rightarrow \infty$, for every possible initial condition x_0. (See Coddington and Levinson [3] for further details.)

As the above discussion suggests, most of the mathematical theory pertaining to systems of the form (2.10) and (2.11) is qualitative in nature. A major computational difference between CVDS and DEDS is that the numerical determination of equilibrium points is significantly simpler than that of calculating invariant probability distributions. Nevertheless, it is our view that the above analogy between CVDS and DEDS is useful in developing an understanding of the major theoretical issues arising from discrete event systems.

III. Generalized Semi-Markov Processes

Consider a discrete event system in which the internal state sequence X has the causal representation (2.7). Then, if $\mathfrak{F}_n = \sigma(X_0, \cdots, X_n)$ is the σ-field corresponding to the history of X up to time n, we find that

$$P\{X_{n+1} \in \cdot \mid \mathfrak{F}_n\} = P(\cdot\,; X_0, \cdots, X_n) \tag{3.1}$$

where the conditional probability appearing on the right-hand side of (3.1) is defined by

$$P(\cdot\,; x_0, \cdots, x_n) = P\{f_{n+1}(\eta_{n+1}, x_0, \cdots, x_n) \in \cdot\}. \tag{3.2}$$

The discrete event system evolves in time by recursively generating X_{n+1} from the conditional distribution (3.2). Once X_{n+1} is obtained, S_{n+1} and Δ_{n+1} can be calculated by using the transformations h_1 and h_2.

Of course, most discrete event systems of interest have more structure than that which we have described above. Specifically, discrete event systems are typically characterized by two different types of entities, namely states and events. In a queueing network, the states generally correspond to queue-length vectors describing the number of customers at each station of the network. The set of events lists the different ways in which the queue-length vector can change as a customer completes service at one station and moves to the next.

To mathematically describe the dynamics of this type of system, we let S denote the (finite or countable) set of states, and E denote the (finite or countable) set of events. A state $s \in S$ is termed a *physical state*, in order to distinguish such states from the state space corresponding to the internal state sequence of the discrete event system. For each $s \in S$, let $E(s)$ be a nonempty finite subset of E denoting the set of events that can trigger transitions out of state s.

Example 1: To model an open queueing network with d stations and one class of customers, we let $S = Z^+ \times Z^+ \times \cdots \times Z^+$ (d times). The vector $s = (s(1), \cdots, s(d)) \in S$ will then represent the queue-lengths (including the customer at the server) at each of the d stations. A state transition occurs via either of the following possibilities: an external arrival event or a departure event. Thus, $E(s) = \{(i, 1): 1 \leq i \leq d\} \cup \{(i, 2): 1 \leq i \leq d, s(i) \geq 1\}$, where $(i, 1)$ corresponds to an external arrival to station i, and $(j, 2)$ denotes a departure from station j.

For each event $e \in E(s)$, we can associate a clock. The reading c_e on clock e can be viewed as representing (in some rough sense) the amount of time that has passed since clock e was last activated.

Example 1 (continued): Suppose that there exists a single server at each of the d stations; each server works at unit rate. For $e = (i, 1)$, c_e corresponds to the amount of time that has elapsed since the last external arrival to station i. For $e = (j, 2)$, c_e denotes the amount of time that has passed since service was initiated on the customer at the server at station j.

The clock readings for event e increase at a speed r_{se}. Thus, the rate at which c_e increases may depend on the physical state occupied by the discrete event system.

Example 1 (continued): If each of the single servers serves at unit rate, then $r_{se} = 1$ for all $s \in S$, $e \in E(s)$. On the other hand, to model a server for which the service rate is proportional to the number of customers in queue at the station, we would set $r_{se} = r_j \cdot s(j)$ for $e = (j, 2)$.

We let $\mathbb{R}^{+}_{-1} = \{-1\} \cup [0, \infty)$ and assume that $c_e \in \mathbb{R}^{+}_{-1}$. We adopt the convention that if $c_e = -1$, then e is currently not active. (For example, event $(j, 2)$ is not active when $s(j) = 0$.) Thus, $C(s) = \{c \in (\mathbb{R}^{+}_{-1})^E: c_e = -1 \text{ iff } e \notin E(s)\}$ is the set of clock readings possible in $s \in S$.

We now give a rough description of the dynamics of the discrete event system. Suppose that at time $\Lambda(n)$, the system has just entered state $s \in S$, and that the clock reading at that instant is represented by $c \in C(s)$. Each clock $e \in E(s)$ will now compete to trigger a transition out of state s. The system evolves by first probabilistically generating, for each event $e \in E(s)$, a "residual lifetime" r.v. which represents the amount of time remaining until event e triggers a transition out of state s. This residual lifetime has a distribution which depends, of course, on the values of c_e and r_{se}. (It may also depend on the history $\mathfrak{F}_n$ in a more complicated way.) Having generated residual lifetimes for all the events $e \in E(s)$, the trigger event e^* is simply the next event to be scheduled by the DEDS. In other words, the trigger event is the event e having the minimum residual lifetime. This minimal residual lifetime therefore yields Δ_{n+1} (i.e., the time between the nth and $(n + 1)$th transitions).

A new physical state s' is now chosen stochastically; its distribution typically depends on both the previous state s and the trigger event e^*. The events $e \in 0(s'; s; e^*) \equiv E(s') - (E(s) - \{e^*\})$ have their clock readings incremented appropriately to reflect the speed r_{se} and the passage of time Δ_{n+1}. (Clocks $e \in 0(s'; s, e^*)$ are "old" clocks that continue to run in state s'.) Clearly, the events $e \in N(s'; s, e^*) \equiv E(s') - 0(s'; s, e^*)$ are the "new" clocks, which will satisfy $c_{s'e} = 0$ at time $\Lambda(n + 1)$. Thus, we have calculated the physical state and clock readings at time $\Lambda(n + 1)$. The process can now be repeated recursively to obtain the physical state/clock readings at $\Lambda(n + 2), \Lambda(n + 3), \cdots$.

We shall now describe the time evolution of the system more precisely. Let $\Sigma = \cup_{s \in S} \{s\} \times [0, \infty) \times C(s)$. The internal state sequence X will be assumed to take values in Σ. Specifically, the internal state at transition n is given by $X_n = (S_n, \Delta_n, C_n) \in \Sigma$. Thus, the first two components of X_n correspond to the physical state and "holding time" of the system at the nth transition. The third component C_n is a vector denoting the state of the clocks at time $\Lambda(n)$. Note that the mappings h_1, h_2 of Section II are given by $h_1(s, t, c) = s$, $h_2(s, t, c) = t$.

If $x_i \in \Sigma$, $i \geq 0$, the vector $\vec{x}_n = (x_0, \cdots, x_n)$ is a possible realization of $\vec{X}_n = (X_0, \cdots, X_n)$. We now need to describe the conditional probability (3.2) (i.e., $P\{X_{n+1} \in \cdot \,|\vec{X}_n = \vec{x}_n\}$) in more detail. To set the stage, we assume that for all $s \in S$, there exists an event $e \in E(s)$ such that $r_{se} > 0$. Thus, in every state, there exists at least one clock with positive speed. We further let $p(s'; \vec{x}_n, e)$ be the (conditional) probability that S_{n+1} equals s', given that $\vec{X}_n$ equals $\vec{x}_n$, and the trigger event e^*_{n+1} at transition $n + 1$ is e. Also, we assume the existence of a family $F(\cdot; \vec{x}_n, e)$ of probability distribution functions such that $F(0; \vec{x}_n, e) = 0$ (i.e., $F(\cdot; \vec{x}_n e)$ corresponds to a positive r.v.). The distribution $F(\cdot; \vec{x}_n, e)$ helps govern the "residual life" of clock e, given that $\vec{X}_n = \vec{x}_n$. We require that for each $\vec{x}_n$, there exists at most one clock e such that $F(t; \vec{x}_n, e)$ is not continuous as a function of t. This guarantees that the trigger event e^*_{n+1} will be uniquely defined for each $\vec{x}_n$. Set

$$F_a(x; \vec{x}_n, e) = F(ax; \vec{x}_n, e), \quad a \geq 0$$

$$\bar{F}_a(x; \vec{x}_n, e) = 1 - F_a(x; \vec{x}_n, e)$$

$$G(dt; \vec{x}_n, e) = F_{r_{s_n,e}}(dt; x_n, e) \prod_{\substack{e' \in E(s_n) \\ e' \neq e}} \bar{F}_{r_{s_n,e'}}(t; \vec{x}_n, e')$$

(s_n is the first component of x_n). Note that $G(dt; \vec{x}_n, e)$ represents the probability, conditional on $\vec{x}_n$, that $\Delta_{n+1} \in dt$, and $e^*_{n+1} = e$. We can now rigorously define the conditional probability structure of the internal state sequence X. For $A = \{s'\} \times [0, T] \times X_{e \in E}(-\infty, a_e]$, let

$$\begin{aligned} P\{X_{n+1} \in A|\vec{X}_n = \vec{x}_n\} \\ = \sum_{e^* \in E(s_n)} p(s'; \vec{x}_n, e^*) \cdot \prod_{e' \notin E(s')} I(a_{e'} \geq -1) \\ \cdot \prod_{e' \in N(s'; s_n, e^*)} I(a_{e'} \geq 0) \cdot \int_{[0,T]} G(dt; \vec{x}_n, e^*) \\ \cdot \prod_{e' \in 0(s'; s_n, e^*)} I(a_{e'} \geq c_{e',n} + t \cdot r_{s_n,e'}) \end{aligned} \tag{3.3}$$

(c_n is the clock reading vector of x_n). The product of indicators over $e' \notin E(s')(e' \in N(s'; s_n; e^*))$ represents the fact that clocks $e \notin E(s')(e' \in N(s'; s_n, e^*))$ have clock readings of $-1(0)$. The product of indicator functions over $0(s'; s_n, e^*)$ corresponds to the fact that the "old" clocks need to be properly incremented to their new values at $\Lambda(n + 1)$.

The conditional probability distribution (3.3) asserts that we may generate X_{n+1} from $\vec{X}_n$ in the following way. We first generate independent r.v.'s $Y_{e,n}(e \in E(S_n))$ from the conditional distributions $F(\cdot; \vec{X}_n, e)$. Then, $\Delta_{n+1} = \min \{Y_{e,n}/r_{S_n,e}: e \in E(S_n)\}$ and the trigger event e^*_{n+1} is the (unique) event $e \in E(S_n)$ which achieves the minimum for Δ_{n+1}. We then generate S_{n+1} from $p(\cdot; \vec{X}, e^*_{n+1})$. Finally, we set

$$\begin{aligned} C_{n+1,e} &= -1, & e &\notin E(S_{n+1}), \\ C_{n+1,e} &= 0, & e &\in N(S_{n+1}; S_n, e^*_{n+1}), \\ C_{n+1,e} &= C_{n,e} + \Delta_{n+1} \cdot r_{S_n,e}, & e &\in 0(S_{n+1}; S_n, e^*_{n+1}). \end{aligned}$$

We call a discrete event system with an internal state sequence X satisfying (3.3) a *(time-inhomogeneous) generalized semi-Markov process*. The term "time-inhomogeneity" reflects the fact that the "residual life" distributions and state transition probabilities $p(s'; \vec{x}_n, e)$ can depend explicitly on the entire history of X. For example, these probability distributions may depend explicitly on $\Lambda(n)$ (i.e., the time of the nth transition); see Section VII for further details. Furthermore, as we shall show in Section V, these processes do indeed extend the notion of the semi-Markov process, thereby justifying use of the term generalized semi-Markov process. For the remainder of this paper, we will refer to discrete-event systems satisfying (3.3) simply as generalized semi-Markov processes.

We conclude this section with an illustration of how we can exploit the causal structure of DEDS satisfying (2.7) to obtain improved statistical efficiency for associated simulations. Specifically, suppose that we wish to calculate, via simulation, an expectation of the form

$$\alpha = Ef(X_0, \cdots, X_n)$$

($f: \Sigma \times \cdots \times \Sigma\ (n + 1 \text{ times}) \to \mathbb{R}$). The standard approach would first replicate the r.v. $f(X_0, \cdots, X_n)$ m independent times, and then form the sample mean of the m observations.

However, an alternative estimator, based on *control variates*, is often available. Suppose there exists a random d-vector χ such that $E\chi = 0$. Such a vector χ is known, in the simulation literature, as a control. Then

$$C(\lambda) = f(X_0, \cdots, X_n) - \lambda^t \chi$$

is an unbiased estimator of α for all $\lambda \in \mathbb{R}^d$. (We adopt the convention that all elements of $\mathbb{R}^d$ are written as column vectors.) Since λ is at our disposal, we may choose λ to minimize var $C(\lambda)$. The optimal value of λ is given by

$$\lambda^* = (E\chi\chi^t)^{-1} \cdot E\chi f(X_0, \cdots, X_n). \quad (3.4)$$

To implement the method of control variates, we generate m independent copies of the pair $(f(X_0, \cdots, X_n), \chi)$. If λ_m is a sample-based estimate for λ^*, we obtain an (asymptotic) improvement over the original estimator by using a sample mean of the $C(\lambda_m)$'s rather than $f(X_0, \cdots, X_n)$'s. The basic idea underlying the use of control variates is that we are "filtering out" the noise in $f(X_0, \cdots, X_n)$ due to χ; this then reduces the variance of the resulting estimator.

The key to the method of control variates is to obtain an easily calculated control χ which is highly correlated with $f(X_0, \cdots, X_n)$. It turns out that the causal structure represented by (2.7) can be used to easily obtain control variates. Suppose that the conditional distribution (2.7) has the property that for some real-valued function g, the conditional mean

$$\tilde{g}(\vec{x}_n) = E\{g(X_{n+1})|\vec{X}_n = \vec{x}_n\}$$

may be easily calculated. For example, if the DEDS is a GSMP, the conditional means $E\{C_{n+1,e}|\vec{X}_n = \vec{x}\}$ often have simple analytical closed-forms.

Let $D_n = g(X_n) - \tilde{g}(\vec{X}_{n-1})$ for $n \geq 1$. If $Eg^2(X_n) < \infty$ for $n \geq 0$, it may be easily verified that the r.v.'s $D_1, D_2, \cdots, D_n$ are martingale differences with respect to the sequence of σ-fields $(\mathfrak{F}_m: m \geq 0)$. Consequently, $D_1, \cdots, D_n$ are orthogonal mean-zero r.v.'s. Since the D_i's have mean zero, it follows that $\chi = (D_1, \cdots, D_n)^t$ is a control. The orthogonality of the D_i's implies that $E\chi\chi^t$ is a diagonal matrix. Thus, (3.4) simplifies to

$$\lambda_i^* = ED_i f(X_0, \cdots, X_n)/ED_i^2 \quad (3.5)$$

provided that $Ef^2(X_0, \cdots, X_n) < \infty$, and $ED_i^2 > 0$. Hence, an advantage of the martingale controls described here is that λ_m need estimate only $2n$ parameters in (3.5), as opposed to $(n^2 + 3n)/2$ in (3.4).

The above discussion shows that the method of control variates is generally applicable to DEDS in which the internal state sequence is causally generated. In particular, martingale control variate schemes can be applied to GSMPs.

IV. Time-Homogeneous GSMPs

In this section, we examine a class of GSMPs, which also satisfy (2.11). As shown in Proposition 2, this will guarantee that X is a time-homogeneous Markov chain.

In many discrete event systems, the constituent conditional probability distributions $F(\cdot; \vec{x}_n, e)$ and $p(\cdot; \vec{x}_n, e)$ defining a GSMP simplify considerably.

Example 1 (continued): Suppose that customers are routed through the queueing network via a substochastic Markovian switching matrix P. Then, the state transition probabilities $p(\cdot; \vec{x}_n, e)$ take the form $p(\cdot; s_n, e)$. This means that the probability distribution of S_{n+1} depends only on the current physical state s_n and the trigger event e. Let $\vec{e}_i$ denote the ith unit vector. The specific form of the state transition probabilities is given by

$$p(s'; s, (i, 1)) = \begin{cases} 1, & \text{if } s' = s + \vec{e}_i \\ 0, & \text{if } s' \neq s + \vec{e}_i \end{cases}$$

$$p(s'; s, (i, 2)) = \begin{cases} P_{ij}, & \text{if } s' = s + \vec{e}_j - \vec{e}_i,\ s(i) \geq 1 \\ 1 - \sum_{j=1}^{d} p_{ij}, & \text{if } s' = s - \vec{e}_i,\ s(i) \geq 1. \end{cases}$$

Suppose that we further assume that the external arrival stream to the ith station is a renewal process with continuous interarrival distribution F_i. Also, suppose that each server employs a first-come/first-serve queueing discipline in which the service requirements for the consecutive customers served at the ith station are i.i.d. with common continuous distribution G_i.

If $e = (i, 1)$, the clock reading $c_{n,e}$ is the amount of time (at the instant $\Lambda(n)$) that has passed since the last customer arrived externally to station i. For $e = (j, 2)$, the clock reading $c_{n,e}$ is interpreted as the amount of service requirement that has been processed on the customer that is in service at station j at time $\Lambda(n)$.

Assume that the service rate of the server at station i is $r_i \cdot s(i)$, so that the rate is proportional to the number of customers in queue at station i. Given the above assumptions, the conditional distribution $F(dt; \vec{x}_n, e) = F(dt; c_{n,e}, e)$ so that the conditional probability distribution function for clock e depends on the history of X only through $c_{n,e}$. For $e = (i, 1)$, the exact form of the conditional distribution is given by

$$\bar{F}(t; c_{n,e}, e) = \bar{F}_i(t + c_{n,e})/\bar{F}_i(c_{n,e})$$

whereas for $e = (j, 2)$,

$$\bar{F}(t; c_{n,e}, e) = \bar{G}_j(t + c_{n,e})/\bar{G}_j(c_{n,e}).$$

Building on the above example, suppose that we have a GSMP for which there exists a family of distributions $(F_e: e \in E)$, and a family of state transition probabilities $(p(\cdot; s, e): s \in S, e \in E(s))$ such that

$$p(\cdot; \vec{x}_n, e) = p(\cdot; s_n, e) \quad (4.1)$$

and

$$\bar{F}(t; \vec{x}_n, e) = \bar{F}_e(t + c_{n,e})/\bar{F}_e(c_{n,e}). \quad (4.2)$$

If a GSMP satisfies the additional conditions (4.1)–(4.2), we refer to the discrete-event system as a *time-homogeneous GSMP*. Noting that $P\{X_{n+1} \in \cdot\, |\vec{X}_n = \vec{x}_n\}$ can be represented in the form $P(x_n, \cdot)$, we see that the internal state sequence of a time-homogeneous GSMP is a time-homogeneous Markov chain.

The Markov structure of a time-homogeneous GSMP can be fruitfully exploited to study the long-run behavior of the corresponding discrete event system. The following result shows that the long-run behavior of the output process $(S(t): t \geq 0)$ can be calculated from that of the internal state sequence X. (In order that Theorem 1 hold, the system need not even be a GSMP.)

Theorem 1: Let f be a real-valued function. Suppose that

$$\frac{1}{n}\sum_{j=0}^{n-1} f(S_j)\Delta_{j+1} \overset{\text{a.s.}}{\to} \mu_1$$

$$\frac{1}{n}\sum_{j=0}^{n-1} \Delta_{j+1} \overset{\text{a.s.}}{\to} \mu_2$$

$$\frac{1}{n}\sum_{j=0}^{n-1} |f(S_j)|\Delta_{j+1} \overset{\text{a.s.}}{\to} \mu_3$$

where μ_1, μ_2, μ_3 are finite r.v.'s with $\mu_2 > 0$ a.s. Then

$$\frac{1}{t}\int_0^t f(S(u))\, du \xrightarrow{\text{a.s.}} \mu_1/\mu_2$$

as $t \to \infty$.

The proof of this result is similar to that of Proposition 2 of Glynn and Iglehart [5] and so is omitted.

To examine the long-run behavior of the Markov chain X, we study the question of existence of invariant probability distributions for X.

Theorem 2: Let X be the Markov chain corresponding to a time-homogeneous GSMP with the following properties:

i) $|S| < \infty$
ii) F_e is continuous with $F_e(0) = 0$, for all $e \in E$
iii) $r_{se} > 0$ for all $s \in S$, $e \in E(s)$
iv) $F_e(c) < 1$ for all $e \in E$, $c < \infty$
v) For all $\epsilon > 0$, $e \in E$, there exists K such that $\bar{F}_e(K + c)/\bar{F}_e(c) < \epsilon$ uniformly in c.

Then, X has an invariant probability distribution π.

For the proof, see the Appendix. Suppose that $F_e(\cdot)$ has a density $f_e(\cdot)$. Then the *hazard rate function* $h_e(\cdot)$ is defined via $h_e(t) = f_e(t)/\bar{F}_e(t)(\bar{F}_e(t) = 1 - F_e(t))$. Condition v) is satisfied if the hazard rate function is bounded below by a positive constant (i.e., inf $\{h_e(t)\colon t \geq 0\} > 0$). It is also satisfied by any finite-mean distribution which is new, better than used, in expectation (see [7] for a discussion of such distributions).

Hypotheses i) and v) of Theorem 2 guarantee that the Markov chain X spends a large fraction of time in compact subsets of Σ; this, in turn, guarantees the type of "positive recurrence" needed to obtain the existence of invariant probability measures. The proof also demands that the transition function of X be continuous in a certain sense; conditions ii)–iv) yield the required continuity.

Let $P_\mu(\cdot)(E_\mu(\cdot))$ denote the probability (expectation) on the path space of X under which X_0 has distribution μ.

Theorem 3: Let X be the Markov chain corresponding to a time-homogeneous GSMP with the following properties:

i) $|S| < \infty$
ii) $\int_{[0,\infty)} tF_e(dt) < \infty$, $e \in E$
iii) For all $s, s' \in S$, $e \in E(s)$ with $p(s'; s, e) > 0$, there exists $e' \in N(s'; s, e)$ such that $r_{s',e'} > 0$.

Then, if X has an invariant probability distribution π, $E_\pi|f(s_0)|\Delta_1 < \infty$.

For the proof of Theorem 3, see the Appendix. The point of Theorem 3 is that it gives sufficient conditions for the finiteness of $E_\pi\Delta_1$ and $E_\pi|f(S_0)|\Delta_1$. Such moment conditions are necessary in order to apply the ergodic theorem. The following result is an immediate consequence of Birkhoff's ergodic theorem and Theorem 1. (We need the continuity of the $F_e(\cdot)$'s in order to guarantee that $P\{\Delta_1 > 0|X_0 = x\} = 1$ for all $x \in \Sigma$. This ensures that $E_\pi\{\Delta_1|\mathcal{G}\} > 0$ a.s.).

Proposition 3: Let X be the Markov chain corresponding to a time-homogeneous GSMP. Suppose that there exists an invariant probability distribution π for X such that $E_\pi\Delta_1 < \infty$ and $E_\pi|f(S_0)|\Delta_1 < \infty$. We further assume that $F_e(\cdot)$ is continuous for all $e \in E$. Then

$$\frac{1}{t}\int_0^t f(S(u))\, du \to \frac{E_\pi\{f(S_0)\Delta_1|\mathcal{G}\}}{E_\pi\{\Delta_1|\mathcal{G}\}} \quad P_\pi \text{ a.s.}$$

as $t \to \infty$, where $\mathcal{G}$ is the invariant σ-field corresponding to X.

Theorems 1–3, together with Proposition 3, give conditions under which a discrete event system "settles down" to a steady-state. It should be emphasized that the results merely assert existence of a steady-state and say nothing about uniqueness. In particular, under the conditions given above, it is quite possible for the system to have multiple steady-state distributions. The particular steady-state distribution governing the discrete event system then depends on the initial state X_0.

Related results on existence of invariant probability measures for time-homogeneous GSMPs appear in König, Matthes, and Nawrotzki [11], [12], and Whitt [14]. The latter paper also gives conditions under which the invariant probability distribution π is continuous in the state-transition probabilities $p(s'; s, e)$ and distributions $(F_e\colon e \in E)$.

In Section VI, we return to this steady-state theme. The regenerative machinery used there establishes both existence and uniqueness results (but under different conditions than those discussed here).

V. Continuous-Time Markov Chains and Semi-Markov Processes

Our objective here is to briefly indicate some of the connections between continuous-time Markov chains, semi-Markov processes, and GSMPs.

Basically, any time-homogeneous GSMP in which all the event distributions $F_e(\cdot)$ are exponential is a continuous time Markov chain. More precisely, consider a time-homogeneous GSMP for which

$$F_e(dt) = \lambda(e)\exp(-\lambda(e)t)\, dt$$

$(\lambda(e) > 0)$ for all $e \in E$. Then, the conditional probability distributions defined by (4.2) take the form

$$\bar{F}(t; \vec{x}_n, e) = \exp(-\lambda(e)t) \tag{5.1}$$

for $t \geq 0$, $e \in E$. Note that the conditional distribution (5.1) is independent of the history $\vec{x}_n$. As a consequence, it is clear that the conditional distribution $G(dt; \vec{x}_n, e)$ depends on $\vec{x}_n$ only through s_n, i.e., $G(dt; \vec{x}_n, e) = G(dt; s_n, e)$. (Recall that in Section IV, $G(dt; \vec{x}_n, e)$ simplified only to the form $G(dt; x_n, e)$.) The specific form of $G(dt; s, e)$ for $s \in S$, $e \in E(s)$, is given by the formula

$$G(dt; s, e) = p(s, e)\, G(dt; s)$$

where

$$G(dt; s) = q(s)\exp(-q(s)t)\, dt$$

$$q(s) = \sum_{e \in E(s)} \lambda(e) \cdot r_{s,e}$$

$$p(s, e) = \lambda(e) \cdot r_{s,e}/q(s).$$

Hence, it follows that

$$\begin{aligned}
P\{S_{n+1} = s', \Delta_{n+1} \in dt|\vec{X}_n = \vec{x}_n\}
&= \sum_{e \in E(s_n)} P\{S_{n+1} = s'|e^*_{n+1} = e, \Delta_{n+1} \in dt, \vec{X}_n = \vec{x}_n\} \\
&\quad \cdot P\{e^*_{n+1} = e, \Delta_{n+1} \in dt|\vec{X}_n = \vec{x}_n\} \\
&= \sum_{e \in E(s_n)} p(s'; s_n, e)\, G(dt; s_n, e) \\
&= p(s'; s_n)\, G(dt; s_n)
\end{aligned} \tag{5.2}$$

where $p(s'; s) = \Sigma_{e \in E(s)} p(s'; s, e) p(s, e)$. It easily follows from (5.2) that $(S_n: n \geq 0)$ is a Markov chain on state space S, and that the Δ_n's are conditionally independent given $(S_n: n \geq 0)$, where

$$P\{\Delta_{n+1} \in dt | S_m: m \geq 0\} = G(dt; S_n).$$

A well-known characterization of continuous time Markov chains then implies that $(S(t): t \geq 0)$ is the (minimal) Markov process corresponding to the generator $Q = (Q(s_1, s_2): s_1, s_2 \in S)$, where

$$Q(s_1, s_2) = p(s_1; s_1)\, q(s_1), \quad s_1 \neq s_2$$
$$Q(s_1, s_2) = -q(s_1), \quad s_1 = s_2.$$

We have therefore calculated the generator of the continuous-time Markov chain associated with a time-homogeneous GSMP in which all clocks are exponential.

Perhaps the most important characteristic of a continuous-time Markov chain is that its long-run behavior may be easily calculated. Specifically, if the internal state sequence X has an invariant probability distribution π, then the "induced" distribution $\bar{\pi}(s) = \pi(\{s\} \times [0, \infty) \times C(s))$ can be determined as the probability solution of the system of linear equations $\bar{\pi}^t R = \bar{\pi}^t$, where $R(s, s') = p(s'; s)$. By contrast, the full set (2.12) of equations for π typically involves solving an integral equation.

We turn now to semi-Markov processes. Consider a time-homogeneous GSMP in which $N(s'; s, e) = E(s')$ for all s', $s \in S$, $e \in E(s)$. In this case, for $n \geq 1$, $C_{n,e} = 0$ for $e \in E(S_n)$, so that

$$F(dt; \vec{X}_n, e) = F_e(dt) \quad \text{a.s.}$$

for $n \geq 1$. As a result, $G(dt; \vec{X}_n, e) = G(dt; S_n, e)$ a.s. for $n \geq 1$, so that $G(dt; \vec{X}_n, e) = P\{\Delta_{n+1} \in dt, e^*_{n+1} = e | \vec{X}_n\}$ depends on the history of X only through S_n. Then, for $n \geq 1$

$$\begin{aligned} P\{\Delta_{n+1} \in dt, S_{n+1} = s' | \vec{X}_n\} &= \sum_{e \in E(S_n)} G(dt; S_n, e)\, p(s'; s_n, e) \\ &= \bar{p}(s', S_n)\, F(dt; S_n, s') \end{aligned} \tag{5.3}$$

where

$$G(dt; s, e) = P\{r_{s,e}^{-1} Y_e \in dt\} \prod_{\substack{e' \in E(s) \\ e' \neq e}} P\{r_{s,e'}^{-1} Y_{e'} > t\}$$

$$(Y_e \text{ has distribution } F_e(\cdot))$$

$$\bar{p}(s'; s) = \sum_{e \in E(s)} G(\infty; s, e)\, p(s'; s, e)$$

$$F(dt; s, s') = \sum_{e \in E(s)} G(dt; s, e)\, p(s'; s, e)/\bar{p}(s'; s).$$

It follows easily from (5.3) that $(S_n: n \geq 0)$ is a Markov chain on state space S (having transition probabilities $\bar{p}(s'; s)$), and that the Δ_n's are conditionally independent r.v.'s given $(S_n: n \geq 0)$, where

$$P\{\Delta_{n+1} \in dt | S_m: m \geq 0\} = F(dt; S_n, S_{n+1}).$$

By definition, $(S(t): t \geq 0)$ is therefore a semi-Markov process. As in the continuous-time Markov chain context, the probability $\bar{\pi}(s) = \pi(\{s\} \times [0, \infty) \times C(s))$ can be calculated as the solution to a suitable system of linear equations (see [2]). Thus, the analytical theory for the steady-state of both continuous-time Markov chains and semi-Markov processes is considerably simpler than that encountered for GSMPs.

We take the view that continuous-time Markov chains and semi-Markov processes play the same role within the DEDS area that linear systems play in the development of CVDS. This is because the full analytical theory of both areas can only be brought to bear in the specialized settings mentioned above.

VI. Regenerative GSMPs

We shall now show that an important class of time-homogeneous GSMPs can be treated as regenerative stochastic processes. Although the results to be described here are far from the most general possible, they are intended to give a flavor of what can be expected in general.

We will consider GSMPs in which the distributions $F_e(\cdot)$ have certain special characteristics. Suppose that a distribution F satisfying $F(0) = 0$ has a density f such that the associated hazard rate $h(t) = f(t)/(1 - F(t))$ is bounded above and below by finite positive constants (i.e., $\inf\{h(t): t \geq 0\} > 0$, $\sup\{h(t): t \geq 0\} < \infty$). The distribution F is then said to be *exponentially bounded*. (We use this term because, for every exponentially bounded distribution, there exist positive constants λ_1, λ_2 such that $\exp(-\lambda_1 t) \leq 1 - F(t) \leq \exp \cdot (-\lambda_2 t)$.)

The tail distribution function of a positive r.v. can be represented in terms of its hazard rate

$$1 - F(t) = \exp\left(-\int_0^t h(s)\, ds\right).$$

Hence, if F is an exponentially bounded distribution with $0 < \alpha \leq h(t) \leq \beta < \infty$,

$$\begin{aligned} F(dt) &= h(t)(1 - F(t))\, dt \\ &\geq \alpha \exp(-\beta t)\, dt \\ &\geq \delta \cdot \beta \exp(-\beta t)\, dt \end{aligned} \tag{6.1}$$

where $\delta = \alpha/\beta$. We will now exploit the above inequality to develop a regenerative structure for GSMPs.

The idea is that under (6.1), we can write

$$F(dt) = \delta \cdot \beta \exp(-\beta t)\, dt + (1 - \delta)\, Q(dt) \tag{6.2}$$

where $Q(dt)$ is a probability distribution function. Thus, $F(dt)$ is, with probability δ, exponential with parameter β. Hence, with positive probability, an exponentially bounded clock acts like a memoryless exponential clock. This, in turn, leads to regeneration.

To be more specific, suppose that $F_e(\cdot)$ is exponentially bounded for all $e \in E$. Let $\alpha(e)$, $\beta(e)$ be the lower and upper bounds on the associated hazard function $h_e(\cdot)$. Then, for each $s \in S$, we have the inequality

$$\begin{aligned} P\{S_1 = s', e_1^* = e, \Delta_1 \in dt | X_0 = (s, t, c)\} \\ \geq \epsilon(s)\, \bar{p}(s'; s, e)\, \bar{q}(s) \exp(\bar{q}(s) t)\, dt \end{aligned} \tag{6.3}$$

where

$$\epsilon(s) = \sum_{e \in E(s)} \alpha(e)/\beta(e)$$

$$\bar{p}(s'; s, e) = p(s'; s, e)\, \beta(e) r_{s,e}/\bar{q}(s)$$

$$\bar{q}(s) = \sum_{e \in E(s)} \beta(e) r_{s,e}.$$

By writing the inequality (6.3) as an equality (in the same fashion as (6.2) was obtained from (6.1)), we see that if $S_0 = s$, then with probability $\epsilon(s)$, (S_1, Δ_1) is independent of X_0. This may appear to suggest that the discrete-event system regenerates with probability $\epsilon(s)$ every time a fixed stated $s \in S$ is hit. Unfortunately, this reasoning is invalid, since (S_2, Δ_2) may still depend on C_0. Therefore, we need to work harder to obtain regeneration.

Assume that the GSMP satisfies:

$$\text{For every } e \in E, \text{ there exists } s \in S \text{ such that } e \in E(s) \text{ and } r_{s,e} > 0. \tag{6.4}$$

$$\text{For every } s, s' \in S, \text{ there exists } e, \bar{s}_1, \bar{e}_1, \cdots, \bar{s}_n, \bar{e}_n \text{ such that} \tag{6.5}$$

$$p(s_1; s, e) r_{s,e} \prod_{i=2}^{n} p(\bar{s}_i; \bar{s}_{i-1}, \bar{e}_{i-1}) r_{\bar{s}_{i-1}, \bar{e}_{i-1}} \cdot p(s'; \bar{s}_n, \bar{e}_n) r_{\bar{s}_n, \bar{e}_n} > 0.$$

Condition (6.5) may be viewed as an irreducibility hypothesis on the GSMP.

Under assumptions (6.4) and (6.5), there exists, for every $s, s' \in S$, a sequence $e, \bar{s}_1, \bar{e}_1, \cdots, \bar{s}_n, \bar{e}_n$ such that the GSMP moves from s to s' with positive probability through the intermediate states $\bar{s}_1, \cdots, \bar{s}_n$, using the successive trigger events $e, \bar{e}_1, \cdots, \bar{e}_n$. In fact, if we set $\bar{s}_0, = s, \bar{s}_{n+1} = s'$, and $\bar{e}_0 = e$, we have the inequality

$$P\{S_i = \bar{s}_i, e_i^* = \bar{e}_{i-1}, \Delta_i \in dt_i, 1 \le i \le n+1 | X_0 = (s, t, c)\} \le \prod_{i=0}^{n} \epsilon(\bar{s}_i) \bar{p}(\bar{s}_{i+1}; \bar{s}_i, \bar{e}_i) \bar{q}(\bar{s}_i) \exp(-\bar{q}(\bar{s}_i) t_{i+1}) \, dt_{i+1}. \tag{6.6}$$

In fact, conditions (6.4) and (6.5) allow us to further choose the path so that each $e \in E(\bar{s}_0)$ appears in the set $\{\bar{e}_0, \cdots, \bar{e}_n\}$. We make this choice of path for the following reason. Note that given $S_0, \cdots, S_{n+1}, e_1^*, \cdots, e_{n+1}^*$, the clock vector C_{n+1} is a function only of $C_0, \Delta_1, \cdots, \Delta_{n+1}$. The right-hand side of (6.6) shows that with probability $\Pi_{i=0}^{n} \epsilon(\bar{s}_i) \bar{p}(\bar{s}_{n+1}; \bar{s}_i, \bar{e}_i)$, the r.v.'s $\Delta_1, \cdots, \Delta_{n+1}$ may be taken as independent exponential r.v.'s. Now, with our choice of path, we can guarantee that for each $e \in E(\bar{s}_{n+1})$, $e \in N(\bar{s}_i; \bar{s}_{i-1}, \bar{e}_{i-1})$ for some $i (1 \le i \le n+1)$. Then, $C_{i,e} = 0$, and it follows that $C_{n+1,e}$ is a function purely of the r.v.'s $\Delta_{i+1}, \cdots, \Delta_{n+1}$ observed along the path. Since the Δ_i's are independent r.v.'s with positive probability, it follows that C_{n+1} is then independent of C_0 (with probability $\Pi_{i=0}^{n} \epsilon(\bar{s}_i) \bar{p}(\bar{s}_{i+1}; \bar{s}_i, \bar{e}_i)$).

For each $s \in S$, let $\bar{\epsilon}(s) = \Pi_{i=0}^{n} \epsilon(\bar{s}_i) \bar{p}(\bar{s}_{i+1}; \bar{s}_i, \bar{e}_i)$ and $L(s)$ be the length of the path constructed above (i.e., the number of states visited). We have shown that if $S_m = s$, then with probability $\bar{\epsilon}(s)$:

1) (S_{m+i}, Δ_{m+i}), $1 \le i \le L(s)$, are independent r.v.'s (S_{m+i} is the ith state on the path constructed for state s, and Δ_{m+i} is a corresponding exponential r.v.), independent of X_m.
2) $C_{m+L(s)}$ is independent of X_m.

1) and 2) imply that $(X_{m+L(s)+k}: k \ge 0)$ is independent of X_m, so that the sequence $((S_{m+L(s)+k}, \Delta_{m+L(s)+k}): k \ge 0)$ is then automatically independent of X_m. Combining this with 1), we conclude that with probability $\bar{\epsilon}(S_m)$, $((S_{m+k}, \Delta_{m+k}): k \ge 1)$ is independent of X_m. Fix a state $s \in S$. We have just shown that there exists a random subsequence $T(1), T(2), \cdots$ of hitting times of s for which the "cycles" $\{((S_j, \Delta_j): T(n) < j \le T(n+1)): n \ge 1\}$ are i.i.d. The random subsequence of regeneration times is obtained from the original hitting time sequence by flipping a coin having probability of success $\bar{\epsilon}(s)$. If the coin flip is successful, then the next $L(s)$ (S, Δ)-tuples are generated using the algorithm described above. This, in turn, gives rise to the desired regenerative structure. We have thus established the following result.

Theorem 4: Consider a time-homogeneous GSMP satisfying (6.4) and (6.5), for which $F_e(\cdot)$ is exponentially bounded for all $e \in E$. If there exists $s \in S$ such that $S_n = s$ infinitely often, $(S(t): t \ge 0)$ is a regenerative process.

An interesting feature of the above regenerative construction is that while the r.v.'s $((S_{m+k}, \Delta_{m+k}): k \ge 1)$ are independent of X_m, it is not true that $(X_{m+k}: k \ge 1)$ is independent of X_m. Thus, while the output process $(S(t): t \ge 0)$ is regenerative, the internal state sequence may not be regenerative. A similar situation arises when we consider the regenerative structure of a continuous-time Markov chain from a GSMP viewpoint. It is well known that the consecutive times at which the chain hits a fixed state constitute regeneration times for the associated (S, Δ) sequence. On the other hand, the full vector C_n of clock readings does not regenerate at such hitting times. In particular, assuming all speeds are unity, the differences between clock readings are preserved from one transition of the full clock sequence to the next. This preservation of memory holds even at transition times to a fixed (physical) state. Thus, the full clock vector does not typically regenerate, even when the GSMP is a continuous-time Markov chain.

Suppose $|S| < \infty$ and fix $s' \in S$. From (6.6), it follows that

$$P\{T(s') > L | X_0 = (s, t, c)\} \ge \epsilon \tag{6.7}$$

for all $(s, t, c) \in \Sigma$, where $\epsilon = \min\{\epsilon(s): s \in S\}$, $L = \max\{L(s): s \in S\}$, $T(s') = \min\{n \ge 1: S_n = s'\}$. A standard "geometric trials" argument then proves that s' is visited infinitely often, yielding the following corollary.

Corollary 1: Consider a time-homogeneous GSMP satisfying (6.4) and (6.5), for which $F_e(\cdot)$ is exponentially bounded for all $e \in E$. If $|S| < \infty$, then $(S(t): t \ge 0)$ is a regenerative process.

A regenerative process is, in some sense, a stochastic process generalization of a sequence of i.i.d. r.v.'s. As a result, we should expect behavior similar to that typical of an i.i.d. sequence; this behavior includes strong laws and central limit theorems.

Theorem 5: Consider a time-homogeneous GSMP satisfying (6.4) and (6.5), for which $F_e(\cdot)$ is exponentially bounded for all $e \in E$. If $|S| < \infty$, there exist finite (deterministic) constants $r(f)$, $\sigma(f)$ such that for every initial distribution μ

$$\frac{1}{t} \int_0^t f(S(u)) \, du \to r(f) \quad P_\mu \text{ a.s.}$$

$$t^{1/2} \left(\frac{1}{t} \int_0^t f(S(u)) \, du - r(f) \right) \Rightarrow \sigma(f) N(0, 1) \quad P_\mu - \text{weakly}$$

as $t \to \infty$.

The proof of Theorem 5 may be found in the Appendix. An important feature of Theorem 5 is that the steady-state limit constants $r(f)$ and $\sigma(f)$ are independent of μ. (Compare this with Proposition 3.) Note also that if we view $f(S(t))$ as the rate at which cost accrues at time t, then the total cost $C(t)$ of running the GSMP up to time t has a distribution which may be approximated (in distribution) as

$$C(t) \approx N(r(f)t, \sigma^2(f)t).$$

The central limit theorem of Theorem 5 also has important implications for steady-state simulations of discrete event systems, since virtually all steady-state confidence interval methodologies (see Chapter 3 of Bratley, Schrage, and Fox, [1]) are based on such a result.

The main point of this section is that regenerative ideas can be applied to discrete event systems. The construction of the associated regeneration times is typically more complicated than that for simpler processes, such as continuous-time Markov chains. Whenever regenerative structure is present, we can expect results similar to Theorem 5. In addition to the regenerative structure identified here, Haas and Shedler [7], [8] have identified regeneration (of a different character) in a number of other GSMP contexts. Thus, we view the laws of large numbers and central limit theorems described here as being typical of a large class of discrete event systems.

VII. Likelihood Ratios for GSMPs

Let $A \subseteq S$, and suppose that we wish to calculate $P\{S(A) \le t\}$, where $S(A) = \inf\{t \ge 0: S(t) \in A\}$. Typically, this probability needs to be numerically calculated; simulation is generally the most popular numerical approach.

In many situations, we expect that $P\{S(A) \le t\}$ is small. For example, if A is the set of "failed states" of a discrete event reliability system, then $P\{S(A) \le t\}$ will be small if the system is reliable. Unfortunately, naive simulation is highly inefficient for such problems; many replications will be necessary in order for the system to experience a reasonable number of failures.

A powerful technique that can be used in such situations is *importance sampling*. The idea is to simulate the system so as to bias it toward failure; the estimator must then be altered so as to compensate for the "biased dynamics." The adjustment factor needed is called a *likelihood ratio*.

Consider a GSMP of the type described in Section III. The probability distributions that govern the dynamics of the system are the conditional distributions $p(\cdot; \vec{x}_n, e)$ and $F(\cdot; \vec{x}_n, e)$. Let $P(\cdot)$ denote the probability distribution of the internal state sequence X under these conditional distributions, and let $E(\cdot)$ be the corresponding expectation operator.

To perform importance sampling, we need to specify the alternative conditional distributions that will appropriately bias the dynamics of the system. For $\vec{x}_n, e$, let $\mathbf{p}(\cdot; \vec{x}_n, e)$ and $\mathbf{F}(\cdot; \vec{x}_n, e)$ be conditional distributions having the property that there exist functions $q(\cdot; \vec{x}_n, e)$ and $f(\cdot; \vec{x}_n, e)$ such that

$$p(\cdot; \vec{x}_n, e) = q(\cdot; \vec{x}_n, e)\, \mathbf{p}(\cdot; \vec{x}_n, e) \qquad (7.1)$$

$$F(dt; \vec{x}_n, e) = f(\cdot; \vec{x}_n, e)\, \mathbf{F}(dt; \vec{x}_n, e). \qquad (7.2)$$

Let $\mathbf{P}(\cdot)$, $\mathbf{E}(\cdot)$ denote the probability distribution and expectation operator corresponding to the conditional distributions $\mathbf{p}(\cdot; \vec{x}_n, e)$ and $\mathbf{F}(\cdot; \vec{x}_n, e)$. The following result is a straightforward generalization of the likelihood ratio ideas in Glynn and Iglehart [6].

Theorem 6: Consider a GSMP with conditional distributions $p(\cdot; \vec{x}_n, e)$, $\mathbf{p}(\cdot; \vec{x}_n, e)$, $F(dt; \vec{x}_n, e)$, $\mathbf{F}(dt; \vec{x}_n, e)$ satisfying (7.1)–(7.2). Let T be a stopping time relative to the internal state sequence X (i.e., $I(T = n)$ is a function of $\vec{X}_n$), and let $Y = f(X_0, \cdots, X_T)$ be real-valued. Then

$$E_\mu YI(T < \infty) = \mathbf{E}_\mu YI(T < \infty)L_T$$

(the equality should be interpreted as: if one expectation exists, then both do, and they are equal), where

$$L_T = \prod_{i=1}^{T} \mathbf{g}(\Delta_i; \vec{X}_{i-1}, e_i^*)\, q(S_i; \vec{X}_{i-1}, e_i^*)$$

and

$$\mathbf{g}(t; \vec{x}_n, e) = f(r_{s,e}t; \vec{x}_n, e) \prod_{\substack{e' \in E(s) \\ e' \ne e}} \cdot\, \bar{F}(r_{s,e'}t; \vec{x}_n, e')/\bar{\mathbf{F}}(r_{s,e'}t; \vec{x}_n, e').$$

Theorem 6 is the key to importance sampling for GSMPs. Rather than replicate copies of the r.v. $YI(T < \infty)$ under P_μ to estimate $\alpha = E_\mu YI(T < \infty)$, we can replicate copies of $YI(T < \infty)L_T$ under $\mathbf{P}_\mu$ to estimate α. By choosing $\mathbf{P}_\mu$ appropriately, significant improvements in computational efficiency over conventional simulation can be achieved.

Likelihood ratio ideas can also be applied to parameter optimization of discrete-event systems. Specifically, likelihood ratio methods can be used to obtain an efficient means of estimating the gradient of the objective function via simulation (see Glynn [4]). This, in turn, can be used to develop a simulation-oriented gradient-based algorithm for optimizing discrete-event processes.

The likelihood ratio methods described here are but two examples of how the GSMP structure of a discrete-event system can be used to obtain computational enhancements to numerical algorithms for discrete-event sytems.

Appendix

Proof of Theorem 2: We first show that X is weakly continuous on the state space Σ, i.e., if $x_n, x \in \Sigma$ and $x_n \to x$ as $n \to \infty$, then $P(x_n, \cdot) \Rightarrow P(x, \cdot)$ as $n \to \infty$, where $\Rightarrow$ denotes weak convergence. (Recall that $P(x, \cdot) = P\{X_{n+1} \in \cdot \mid X_n = x\}$ is the transition function of X.)

Let $(V(e, c_e): e \in E)$ be a collection of independent r.v.'s having marginal distributions specified by

$$P\{V(e, c_e) > t\} = \bar{F}_e(t + c_e)/\bar{F}_e(c_e).$$

Then, under (4.1) and (4.2), the conditional probability distribution $G(\cdot; \vec{x}_n, e) = G(\cdot; x_n, e)$, where

$$G(u; (s, t, c), e) = Ef(V(e', c_{e'}): e' \in E(s))$$

and

$$f(v(e'): e' \in E(s)) = I(\min_{e' \in E(s)} v(e')/r_{s,e'} \le u,\ \operatorname{argmin}_{e' \in E(s)} v(e')/r_{s,e'} = e).$$

By conditions ii) and iii), it is evident that $\bar{F}_e(t + c)/\bar{F}_e(c)$ is continuous in c at every t. Hence, $(V(e, c_e'): e \in E) \Rightarrow (V(e, c_e): e \in E)$ whenever $(c_e': e \in E) \to (c_e: e \in E)$. Since f is continuous at $(V(e, c_e): e \in E)$ (we use ii)–iv) here), it follows that $G(\cdot; x', e) \Rightarrow G(\cdot; x, e)$ whenever $x' \to x$. Thus, the distributions of the trigger event e^* and Δ are (weakly) continuous in $x = (s, t, c)$. Consequently, the distribution of $X_1 = (S_1, \Delta, C_1)$ is (weakly) continuous in $X_0 = (s, t, c)$, thereby proving the required continuity of X.

Fix $x \in S$. The second step involves showing that $\{\mu_n: n \geq 1\}$ is tight, where

$$\mu_n(\cdot) = \frac{1}{n} \sum_{j=0}^{n-1} P\{X_j \in \cdot \,|X_0 = x\}.$$

To establish tightness, we use the Lyapunov function

$$g(x) = t + \sum_{e \in E} |c_e|;$$

condition v) guarantees that $E\{g(X_1)|X_0 = x\} \leq g(x) - \epsilon (\epsilon > 0)$ uniformly in x outside a compact set, and this yields tightness.

Prohorov's theorem then asserts the existence of a subsequence n' and a probability π on Σ such that $\mu_{n'} \Rightarrow \pi$. A standard argument (see, for example, [10]) then uses the weak continuity of X to prove that π is, in fact, invariant for X.

Proof of Theorem 3: We first note that since $|S| < \infty$, f is bounded, and it suffices to prove only that $E_\pi \Delta_1 < \infty$. By the stationarity of X under π, this is equivalent to showing that $E_\pi \Delta_2 < \infty$. This will follow if we can prove that $E\{\Delta_2|X_1 = x\}$ is uniformly bounded in x.

For s, s', $e \in E(s)$, let

$$M(s'; s, e) = \min_{e' \in N(s'; s, e)} V(e', 0)/r_{s', e'},$$

$$m(s'; s, e) = EM(s'; s, e).$$

Clearly, $m(s'; s, e) \leq \Sigma_{e \in E} \int_{[0,\infty)} tF_e(dt)/r$, where r is the minimum over the speeds $r_{s',e'}$ of condition iii). To finish the proof, note that $C_{1,e} = 0$ for $e \in N(S_1; S_0, e_1^*)$, and hence $E\{\Delta_2|X_1 = x\} \leq Em(S_1; S_0, e_1^*) \leq \Sigma_{e \in E} \int_{[0,\infty)} tF_e(dt)/r$.

Proof of Theorem 5: We need to verify the hypotheses of the regenerative strong law of large numbers and central limit theorem (see Smith [13]). Fix $s' \in S$, and let τ be the first $m \geq L(s')$ such that the next $L(s')$ states to be visited (after m) form the specified path for s'. (We will base our regenerations on visits to s'.)

Since $|S| < \infty$, f is bounded so that it suffices to verify that the moment

$$E\left\{\left(\sum_{i=1}^{\tau} \Delta_i\right)^2 \middle| X_0 = x\right\}$$

is bounded in x. We first note that there exists L such that

$$\sup_x P\{\tau > L|X_0 = x\} < 1$$

(use (6.7)). Since the tail of τ is then geometrically dominated, this implies that $E\{\tau^p|X_0 = x]$ is bounded in x, for all $p > 0$. Now

$$\begin{aligned} E&\left\{\left(\sum_{i=1}^{\tau} \Delta_i\right)^2 \middle| X_0 = x\right\} \\ &\leq E\{\tau^2 \max_{1 \leq i \leq \tau} \Delta_i^2|X_0 = x\} \\ &\leq E^{1/2}\{\tau^4|X_0 = x\}\, E^{1/2}\{\max_{1 \leq i \leq \tau} \Delta_i^4|X_0 = x\} \\ &\leq E^{1/2}\{\tau^4|X_0 = x\}\, E^{1/2}\left\{\sum_{i=1}^{\tau} \Delta_i^4|X_0 = x\right\}, \end{aligned}$$

so we need only show the second factor above is bounded in x. But

$$\begin{aligned} E\left\{\sum_{i=1}^{\tau} \Delta_i^4|X_0 = x\right\} &= \sum_{i=1}^{\infty} E\{\Delta_i^4 I(\tau \geq i)|X_0 = x\} \\ &\leq \sum_{i=1}^{\infty} E^{1/2}\{\Delta_i^8|X_0 = x\} \\ &\quad \cdot P^{1/2}\{\tau \geq i|X_0 = x\}. \end{aligned}$$

Again, since the tail of τ is geometrically dominated, it suffices to prove that $E\{\Delta_i^8|X_0 = x\}$ is bounded in x. But for $x = (s, t, c)$, we can select $e \in E(s)$ so that $r_{s,e} > 0$. Then

$$\begin{aligned} E\{\Delta_1^8|X_0 = x\} &\leq \int_0^\infty \frac{(1 - F_e(c_e + r_{s,e} t^{1/8}))\, dt}{1 - F_e(c_e)} \\ &= \int_0^\infty \exp\left(-\int_{c_e}^{c_e + r_{s,e} t^{1/8}} h_e(u)\, du\right) dt \\ &\leq \int_0^\infty \exp\left(-\alpha(e) r_{s,e} t^{1/8}\right) dt. \end{aligned}$$

Since $|S| < \infty$, it is evident that $E\{\Delta_1^8|X_0 = x\}$ is bounded in x, from which it follows that $E\{\Delta_i^8|X_0 = x\}$ is bounded in x.

References

[1] P. Bratley, L. Schrage, and B. L. Fox, *A Guide to Simulation.* New York, NY: Springer-Verlag, 1987.

[2] E. Cinlar, *An Introduction to Stochastic Processes.* Englewood Cliffs, NJ: Prentice-Hall, 1975.

[3] E. A. Coddington and N. Levinson, *Theory of Ordinary Differential Equations.* New York, NY: McGraw-Hill, 1955.

[4] P. W. Glynn, "Likelihood ratio gradient estimation: an overview," in *Proc. of the 1987 Winter Simulation Conference*, pp. 366–375, 1987.

[5] P. W. Glynn and D. L. Iglehart, "Simulation methods for queues: an overview," in *Queueing Systems: Theory and Applications*, pp. 221–256, 1988.

[6] ——, "Importance sampling for stochastic simulations," *Management Sci.*, to appear 1989.

[7] P. J. Haas and G. S. Shedler, "Recurrence and regeneration in non-Markovian networks of queues," *Stochastic Models*, vol. 3, pp. 29–52, 1987.

[8] ——, "Regenerative generalized semi-Markov processes," *Stochastic Models*, vol. 3, pp. 409–438, 1987.

[9] Y. C. Ho, "Performance evaluation and perturbation analysis of discrete event dynamical systems," *IEEE Trans. Automat. Contr.*, vol. 32, pp. 563–572, 1987.

[10] A. F. Karr, "Weak convergence of a sequence of Markov chains," *Z. Wahrscheinlichkeitstheorie verw. Gebiete*, vol. 33, pp. 41–48, 1975.

[11] D. König, K. Matthes, and K. Nawrotzki, *Verallgeneinerungen der Erlangschen und Engsetschen Formeln.* Berlin, West Germany: Akademie-Verlag, 1967.

[12] ——, "Unempfindlichkeitse Igenschaften Von Bedien-Ungsprozessen." Appendix to B.V. Gnedenko and I. N. Kovalenko, *Introduction to Queueing Theory.* Berlin: West Germany: Akademie-Verlag, 1974.

[13] W. L. Smith, "Regenerative stochastic processes," *Proc. Roy. Soc. Ser. A*, vol. 232, pp. 6–31, 1955.

[14] W. Whitt, "Continuity of generalized semi-Markov processes," *Math. Oper. Res.*, vol. 5, pp. 494–501, 1980.

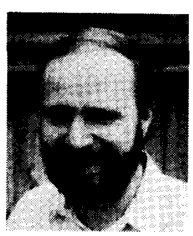

Peter W. Glynn received the Ph.D. degree from Stanford University, Stanford, CA, in 1982.

He spent five years as an Assistant Professor in the Department of Industrial Engineering at the University of Wisconsin-Madison. He is currently an Associate Professor in the Department of Operations Research at Stanford University. His research interests include stochastic systems, computational probability, and simulation.