Guest Editors' Introduction to Special Issue Honoring Donald L. Iglehart

This special issue was conceived to honor Donald L. Iglehart on the occasion of his 80th birthday. He has profoundly influenced the simulation community throughout the course of his career. By connecting the design and analysis of simulation-based estimation algorithms to the underlying theory of stochastic processes, he introduced a new level of rigor to the simulation field. To quote his citation for the 2012 INFORMS Simulation Society Lifetime Professional Achievement Award: "It is no exaggeration to say that Don Iglehart's contributions made simulation a respectable research discipline in some circles of the Operations Research community." His work on simulation—as well as on inventory theory and on diffusion limits and approximations for heavily congested stochastic systems—earned him the John von Neumann Theory Prize in 2002. He is also a member of the National Academy of Engineering and a Fellow of both INFORMS and the Institute of Mathematical Statistics.

In a series of papers beginning in 1975, Iglehart introduced and led the development of the regenerative method for simulation output analysis, providing the first sound basis for computing point and interval estimates for measures of steady-state system performance. A large body of subsequent work by Iglehart, his students, and his colleagues both deepened and extended the underlying theory—for example, it has been shown that any "well posed" steady-state simulation has a type of regenerative structure—and also refined the method in the setting of specific classes of stochastic models such as networks of queues, stochastic Petri nets, and generalized semi-Markov processes. Even for simulations in which regenerative simulation per se may be infeasible, regenerative arguments can be used to establish strong laws of large numbers and central limit theorems, showing that a given steady-state estimation problem is well defined, and hence perhaps amenable to other types of output analysis methods. In other seminal work, Iglehart established a general theoretical framework for the study of standardized time series methods, which encompass many well-known output analysis algorithms such as batch means and spectral methods. Don has been an inspiring teacher, mentor, and colleague. The six articles in this special issue were solicited from his academic progeny and attest to his continuing impact on simulation research.

Initial work on regenerative simulation focused on a single sequence of regeneration points at which the underlying stochastic process of a simulation probabilistically restarts. A classic example is given by the successive times at which there is an arrival to an empty GI/G/1 queue: at each regeneration time, the future evolution of the number-in-system process is probabilistically the same and is independent of the history of the process prior to regeneration. Among other things, this implies that the steady-state distribution of the process is determined by its behavior within a regenerative "cycle" (an interval demarcated by a pair of successive regeneration points). For other stochastic processes, such as positive recurrent Markov chains on a discrete state space, there are multiple sequences of regeneration points. In the article "Resampled Regenerative Estimators" by James Calvin and Marvin Nakayama, the authors compare new and existing methods that exploit the multiple sequences to reduce the variance of simulation-based estimates of steady-state performance measures relative to the standard regenerative method. All of the methods considered reduce the variance but can have different bias properties.

DOI: http://dx.doi.org/10.1145/2822375

A difficulty in applying the regenerative method is that it can be hard to identify a sequence of regeneration points for a complex stochastic system. For example, in a large network of queues, regeneration point sequences analogous to the GI/G/1 queue example are unlikely to exist—too many stochastic activities are going on simultaneously at each time point. One tactic for ameliorating this problem is to employ more subtle regenerative sequences. Specifically, suppose that we can represent each service-time or interarrival-time distribution as a mixture of distributions in which one of the mixture components is an exponential distribution. When generating a service or interarrival time during the simulation, one conceptually flips a coin to select the component from which to generate. At certain time points when all scheduled service and interarrival times have been generated from their exponential components, the occurrence of a regeneration can be deduced from the memoryless property of the exponential distribution. Such mixture representations have been established for light-tailed distributions having bounded hazard rates. The article “Regenerative Simulation for Queueing Networks with Exponential or Heavier Tail Arrival Distributions” by Sarat Moka and Sandeep Juneja extends the foregoing approach to a broad class of queueing networks. A key innovation is the representation of an interarrival time as a sequence of phases, at least one of which has a positive probability of being generated from an exponential distribution. This technique is potentially applicable to a broad range of stochastic systems beyond queuing networks.

An important contribution of Iglehart and colleagues was to recognize the potential of the generalized semi-Markov process (GSMP) as a unifying framework for studying foundational issues in discrete-event simulation. The GSMP model incorporates the fundamental building blocks of “events” and “states” as well as “clocks” that measure the times until events are scheduled to occur. This work gave rise to a line of research aiming to identify conditions on the building blocks of a GSMP model under which steady-state estimation problems are well defined and simulation methods such as the regenerative method are applicable. Usually, the key challenge is to show that the GSMP is “recurrent” in that, with probability 1, it visits every state infinitely often. The article “On Transience and Recurrence in Irreducible Finite-State Stochastic Systems” by Peter W. Glynn and Peter J. Haas highlights the rich and complex behavior of the GSMP model by showing that, in marked contrast to Markov and semi-Markov processes, irreducible finite-state GSMPs with heavy-tailed clock-setting distributions can have transient states. The authors provide a rough analogy between recurrence theory for GSMPs and the famous Chung-Fuchs results for random walks on a d-dimensional integer lattice.

One indirect way in which Iglehart has influenced simulation methodology derives from his work on diffusion limits and approximations for queues with highly utilized servers. A 1989 paper by Ward Whitt showed that such results can be used to plan the length and number of queueing simulations even before any data has been collected. In this setting, a common belief is that long simulation run lengths are required to accurately estimate certain steady-state performance measures when servers are heavily utilized. The article “How Hard Are Steady State Queueing Simulations?” by Erik Ni and Shane Henderson invites the reader to rethink this assumption. Typically in practice, the length of a queue is held in check because customers abandon the queue if they have to wait too long or because the queue is operated in a “quality and efficiency” regime where servers are heavily utilized but wait times are short. Ni and Henderson show that moderate run lengths can suffice for such “well dimensioned” systems.

The article “Exact Sampling of Stationary and Time-Reversed Queues” by Jose Blanchet and Aya Wallwater also reflects Iglehart’s twin interests in queueing theory and steady-state simulation. The authors provide a novel algorithm for exact simulation of the stationary waiting time sequence of a single server queue backward in time, and show that their method is valid under minimal assumptions. Such simulation is a key ingredient in Dominated Coupling from the Past (DCFTP) methods. DCFTP allows exact generation of the steady-state distribution of a Markov chain with unbounded state space. (Thus, rather than using special estimation methods on an autocorrelated nonstationary simulation, as in the regenerative and standardized time-series methods, one can use ordinary estimation methods by simulating independent and identically distributed copies of the stationary version.) A dominating stationary process is simulated backward in time until it coincides with the Markov chain of interest; the sample path of the chain is then reconstructed forward to time zero, at which point the state of the chain is a sample from the target stationary distribution. The single-server queue process studied in the article can be used as a dominating process in a broad variety of DCFTP settings.

A basic step in simulating a queueing or other service system is to first specify the arrival process. When there is available data composing a sequence of observed interarrival times, Kolmogorov-Smirnov (KS) tests are often applied to confirm or reject a hypothesis that these times are independent and identically distributed (i.i.d.) samples from a specified continuous distribution function. Such testing can also be used for simulation output analysis. For a specified significance level, the power of the test can be increased by transforming the data prior to testing. This issue is explored in the article "The Power of Alternative Kolmogorov-Smirnov Tests Based on Transformations of Data" by Song-Hee Kim and Ward Whitt. The authors propose a new procedure wherein the observations are first transformed to (putatively i.i.d.) unit exponential random variables, and hence can be viewed as interarrival times in a Poisson process. Then the transformed data are subjected to a test for Poisson processes due to Peter Lewis. The authors show that other KS variants essentially focus on cumulative sums of the interarrival times, which leads to lower discriminative power.

We thank the contributing authors for their hard work in writing and revising their articles, as well as for acting as referees. We also thank James R. Wilson for his refereeing efforts and for providing overall editorial coordination. Finally, we thank TOMACS Editor-in-Chief Adelinde Uhrmacher for her support throughout.

Peter W. Glynn and Peter J. Haas
Stanford University and IBM Research
Guest Editors