TUTORIAL: TEACHING AN ADVANCED SIMULATION TOPIC

Shane G. Henderson

School of Operations Research and Information Engineering Cornell University Ithaca, NY 14850, USA

Sheldon Jacobson

Stewart Robinson

Department of Computer Science University of Illinois Urbana, IL 61801, USA School of Business and Economics Loughborough University Loughborough, Leicestershire, UK

ABSTRACT

In this tutorial three experienced instructors describe how they teach an advanced simulation topic: Poisson processes (Henderson), variance reduction (Jacobson), and verification and validation (Robinson).

1 HENDERSON: POISSON PROCESSES

1.1 Definitions

In the simulation community we tend to think of Poisson processes as being a convenient model for representing customer arrival times in some service system. Accordingly, we usually define a Poisson process as a counting process on the real line with additional properties. One way to define Poisson processes uses "little o" notation as in Ross (1996), Definition 2.4.1. A slightly more convenient form is as in Ross (1996), Definition 2.1.1 for homogeneous Poisson processes. An adaptation of that definition for nonhomogeneous Poisson processes is as follows.

Definition 1 A process $(N(t): t \ge 0)$ is a *counting process* if $N(0) \ge 0$, N(t) is integer valued for all t and non-decreasing. A counting process N is a Poisson process with rate function $(\lambda(t): t \ge 0)$ if N(0) = 0 and

- 1. The number of events N(t) N(s) in the interval (s,t] is Poisson distributed with mean $\int_{s}^{t} \lambda(u) du$.
- 2. The number of events in non-overlapping intervals are independent random variables.

I have taught Poisson processes in my Ph.D. simulation course from this definition, and found that it works, but that proofs of results on thinning, composition and transformation are somewhat clumsy. In particular, I have found that proving independence in results on thinning is cumbersome. I have also found this approach to be a bit clumsy when I want to model call arrivals in time *and space* as in applications in the emergency services.

There is another view of Poisson processes that is more general, and in which the proofs of the above results are more natural. It is based on the notion of Poisson *point processes*. An excellent reference that communicates the right ideas at the right level is Resnick (1992), Chapter 4, and I have taken much of the material that follows from there. A fully measure-theoretic and less accessible treatment can be found in Cinlar (2011), Chapter VI). In this section I explain what Poisson point processes are, and demonstrate how various results that are important in simulation can be proved both intuitively and with relative ease

using this framework. It is not clear to me that one can do *more* with this more general definition than is possible in the more familiar framework above, but certainly the results seem more elegant and intuitive to me when derived in this framework.

A point process is, intuitively speaking, a collection of random points. More formally, let *E* be a subset of \mathbb{R}^d , and let $(X_n : n \ge 1)$ be a collection of random elements of *E*. We think of these elements as the *points*. For a set $A \subset E$ (it needs to be measurable, but we suppress measurability issues here), let

$$N(A) = \sum_{n} I(X_n \in A)$$

be the number of points that fall in the set *A*, where $I(\cdot)$ is the indicator function that is 1 if its argument is true, and 0 otherwise. The quantity N(A) is random, since the elements $(X_n : n \ge 1)$ are. Set $\mu(A) = EN(A)$ to be the expected number of points that fall in the set *A*. We then say that $N(\cdot)$ is a point process with mean measure μ .

Definition 2 A *Poisson point process* is a point process with the additional properties below.

1. For $A \subset E$ and $k \ge 0$,

$$P(N(A) = k) = \begin{cases} \frac{e^{-\mu(A)}\mu(A)^k}{k!} & \mu(A) < \infty\\ 0 & \mu(A) = \infty. \end{cases}$$

2. If A_1, A_2, \ldots, A_m are disjoint subsets of *E*, then $N(A_1), N(A_2), \ldots, N(A_m)$ are independent random variables.

Definitions 1 and 2 are very similar. In Definition 1 N(t) is the number of events in the interval [0,t] which, in the new notation, is N([0,t]) (and we take $E = [0,\infty)$). The two conditions in the two definitions have parallel meanings, although Definition 2 is more general because the points can be in a higher dimension.

As a first example, the superposition property of independent Poisson point processes follows essentially immediately from these definitions. In particular, if one merges the points of two independent Poisson point processes on the same set E, then another Poisson point process results, as can be checked by verifying Properties 1 and 2 of Definition 2.

1.2 Transformations and Inversion

Our first result proves that deterministically transformed points from a Poisson process yield a new Poisson process. This result can be used, for example, to prove the validity of the "inversion" method (Cinlar 1975) for generating nonhomogeneous Poisson processes from homogeneous Poisson processes on the real line.

Proposition 1 Let *N* be a Poisson point process with points $(X_n : n \ge 1)$ on $E \subset \mathbb{R}^d$. Let $f : E \to E'$ be a function mapping *E* to $E' \subset \mathbb{R}^{d'}$. Let N'(A') be the number of transformed points $(f(X_n) : n \ge 1)$ lying in the set $A' \subset E'$. Then $N'(\cdot)$ is a Poisson point process on E' with mean measure $\mu' = \mu \circ f^{-1}$.

Proof. We need to verify the two properties of a Poisson point process. For a fixed $A' \subset E'$, the number of transformed points N'(A') in A' is the number of points in the pre-image $f^{-1}(A')$, where $f^{-1}(A') = \{x \in E : f(x) \in A'\}$. But this is just $N(f^{-1}(A'))$ which is Poisson distributed with mean $\mu(f^{-1}(A'))$ as required. As to the second property, if A'_1, A'_2, \ldots, A'_m are disjoint in E', then

$$A_1 = f^{-1}(A'_1), A_2 = f^{-1}(A'_2), \dots, A_m = f^{-1}(A'_m)$$

are disjoint in *E*, and therefore $(N'(A'_k) = N(A_k) : k = 1, 2, ..., m)$ are independent as they inherit that same property from *N*.

This result establishes the validity of the inversion method for generating nonhomogeneous Poisson processes. Indeed, suppose that $\Lambda : [0, \infty) \to [0, \infty)$ is the desired cumulative intensity function of a (nonhomogeneous) Poisson process on $[0, \infty)$, and suppose that (as the inversion method we are familiar with requires) the *i*th event time for such a process is constructed as $\Lambda^{-1}(X_i)$, where $(X_i : i \ge 1)$ are the event times in a Poisson process with constant rate function 1. To see why this would yield the desired process using Poisson point processes, start with a homogeneous Poisson point process with mean measure defined by $\mu((0,t]) = t$, and take $f = \Lambda^{-1}$ in Proposition 1 above. The transformed points are then a Poisson point process on $[0,\infty)$ with mean measure defined by $\mu'((0,t]) = \mu(f^{-1}((0,t])) = \mu((\Lambda(0),\Lambda(t)]) = \Lambda(t) - \Lambda(0)$ as required.

Proposition 1 was also used to establish the validity of a "projection" technique for generating Poisson processes in Saltzman et al. (2012). We will see an example using this technique in Section 1.4 below.

1.3 Marking and Thinning

The second major technique for generating Poisson processes on the line after inversion is thinning (Lewis and Shedler 1979). A clean derivation of thinning can be based on Poisson point processes through marking, thinning and transformations. A marking of a point process involves associating additional random variables to each point in the process. Markings can be very general, but for our purposes we content ourselves with a fairly specific marking. Perhaps surprisingly, when we mark the points of a Poisson point process with i.i.d. marks, we get another Poisson point process. This is not the most general marking that still yields another Poisson point process, but it is sufficient for our purposes.

Proposition 2 Suppose $(X_n : n \ge 1)$ are the points of a Poisson point process on $E_1 \subset \mathbb{R}^{d_1}$ with mean measure μ_1 . Suppose $(Y_n : n \ge 1)$ are i.i.d. random vectors taking values in $E_2 \subset \mathbb{R}^{d_2}$ with measure μ_2 (so that $P(Y_1 \in A_2) = \mu_2(A_2)$). Then the points $(Z_n : n \ge 1)$ with $Z_n = (X_n, Y_n)$ are those of a new Poisson point process on $E = E_1 \times E_2$ with mean measure $\mu = \mu_1 \times \mu_2$, so that if $A = A_1 \times A_2$ with $A_1 \subset E_1$, $A_2 \subset E_2$, then $\mu(A) = \mu_1(A_1)\mu_2(A_2)$.

For a proof, see Resnick (1992), Proposition 4.10.1.

We can now use this marking idea to derive the thinning technique for Poisson processes. We state this result without specializing to processes on the real line, because it is more general than that setting.

Proposition 3 Suppose that $(X_n : n \ge 1)$ are the points of a (nonhomogeneous) Poisson point process N with mean measure μ on $E \subset \mathbb{R}^d$. Suppose that $\mu(A) = \int_A \lambda(u) du$, so that $\lambda(\cdot)$ is the usual rate function (also called the local intensity) of N. If we accept a point at x with probability p(x) independently of all other points, then the process of accepted points is a Poisson point process with rate function $(\lambda(x)p(x) : x \in E)$, and is independent of the Poisson point process of rejected points.

Proof. Mark the points $(X_n : n \ge 1)$ with i.i.d. uniform random variables $(U_n : n \ge 1)$ to yield the points $((X_n, U_n) : n \ge 1)$ of a Poisson point process in \mathbb{R}^2 with mean measure $\mu \times \nu$, where ν is the distribution of a uniform random variable on (0, 1). Let

$$((Y_k, V_k) = (X_{n(k)}, U_{n(k)}) : k \ge 1)$$

be the subset of points that lie under the boundary $((x, p(x)) : x \in E)$, i.e., those points that satisfy $U_{n(k)} \leq p(X_{n(k)})$. These points are independent of the rejected points that lie above the boundary, because Poisson point processes are independent in disjoint regions. They form a Poisson point process with the same mean measure $\mu \times v$, but on the restricted domain $E' = \{(x, u) : x \in E, u \leq p(x)\}$. Now apply the projection map f(x, u) = x to these points to obtain the locations associated with accepted points. By Proposition 1 these points are those of a Poisson point process with mean measure defined by

$$\mu \times v(\{(x,u) : x \in A, u \le p(x)\}) = \int_A \lambda(x)p(x) dx$$

as required. The rejected points are those that lie above the surface, and similarly form another Poisson process with rate function $\lambda(\cdot)(1-p(\cdot))$.

Proposition 3 shows that thinning a (nonhomogeneous) Poisson process yields another Poisson process and it holds in any dimension. It is easily extended to allow one to classify points into one of several categories (not just two categories), independently at each point, according to probabilities that depend on the location of the point. Independence of the categories of points from one another is a consequence of the fact that Poisson point processes are independent in disjoint sets.

Marking can be useful in other contexts as well. For example, consider modeling the arrival process of calls to an ambulance service. The times of the calls $(T_n : n \ge 1)$ might be modeled as a Poisson process on $[0,\infty)$. The locations of calls $((X_n, Y_n) : n \ge 1)$ might be (conditionally) generated from a distribution that depends on the time of the call. From Proposition 2, we can also view the generated times and call locations jointly as a Poisson process in three dimensions. This may yield more efficient generating algorithms than generating the times and then marking them as described above. Also, thinning allows us to conclude, e.g., that the process of calls that appear in a particular subdivision is Poisson. However, we cannot conclude that the process of calls for which ambulances arrive within a given time threshold is Poisson, since the "thinning" is not independent from point to point.

1.4 A Last Example

To see how the above results may be applied in another example, suppose one wants to generate a homogeneous Poisson process N over the unit circle, represented in polar coordinates as $\{(r, \theta) : 0 \le r \le 1\}$. Extracting the value of r from each point in such a process is simply a projection, and so by Proposition 1, the set of r values so obtained forms a Poisson point process that we call R. To generate R, we generate the Poisson point process on [0,1] with mean measure defined via $\mu((0,x]) = \lambda \pi x^2$, where λ is a constant giving the desired intensity of N. To do so we can use inversion or thinning. Then, we simply mark each point with i.i.d. uniform angles θ on the interval $(0, 2\pi]$ to obtain N.

If we instead wanted to generate a homogeneous Poisson process over an elliptical region, then we can simply generate a homogeneous Poisson process over the circle as above, and then transform those points into an ellipse, exactly as described in Rubinstein and Melamed (1998), p. 50.

2 JACOBSON: VARIANCE REDUCTION TECHNIQUES

Variance reduction techniques (VRTs) exploit the controlled structure inherent in discrete event simulation models. Through a variety of input and/or output manipulations, simulation output performance measure estimators can be designed that have lower variance than standard or crude average estimators. The use of variance reduction techniques creates an environment whereby something is gained (i.e., a more accurate estimator) at a minimal cost (i.e., little or no data are collected), resulting in an improved use of simulated data and better decision-making.

When teaching variance reduction techniques, the motivating factor is that such techniques have the potential to save time and effort when testing hypotheses and accurately estimating simulation performance measures, based on the output data collected from a simulation model. A simple analogy to variance reduction techniques is risk reduction in portfolio design and management. Given a set of investment options, each with their own expected return and risk, the objective is to design a portfolio that maximizes the overall total return, while minimizing investor risk exposure. For simulation outputs, the expected return corresponds to the expected value of a simulation performance measure estimator, while the risk exposure corresponds to the variance of the estimator. The manner in which a portfolio is assembled is analogous to how the estimator is designed in using simulation output data to achieve an accurate and precise (i.e., unbiased with minimal variance) estimator. Given that variance reduction techniques come in numerous forms, each specific technique is motivated by how the simulation analyst exploits the power of controlling the randomness that drives a simulation or is outputted from a simulation model.

At the very foundation of variance reduction techniques is controlled data dependency. This is often a new concept for students, since they are often exposed to classical statistics that routinely assumes (requires) data independence. Given the control that simulationists have with their models, students are challenged to rethink their view of data dependency as their ally rather than their adversary.

Common and antithetic random numbers are a simple and valuable tool to induce dependency between simulation runs. Common random numbers use identical random number streams to drive two or more simulation runs, while antithetic random numbers use identical random number streams that are flipped (i.e., a U(0,1) random variable U is replaced by 1-U) to drive two or more simulation runs. Using a simple single server queue, one can explain the relationship between how such number streams induce such dependencies, and how the resulting output dependencies can be best exploited to design lower variance estimators. The key concept to communicate is that when executing a simulation model twice, if identical random number streams are used to drive the model runs, the output streams will be identical. By analyzing the trends that are likely to occur between the simulation inputs and the simulation outputs, one can use common random number streams for some input streams, and antithetic random number streams for other input streams so as to induce a predictable dependency between the output streams. To illustrate this point, in a single server queue, long interarrival times often means that customer arrival spacing is sufficiently large that waiting times in the queue tend to be smaller. Similarly, short interarrival times often means that customer arrival spacing is sufficiently short that waiting times in the queue tend to be larger. Simple reasoning as such makes it possible to induce negative dependency between performance measures that are directly related to the queue waiting times, for two simulation runs, using antithetic random number streams for the interarrival input stream. Similar (albeit reversed) reasoning can be applied to the customer service inputs streams. The key point to communicate to students is that input stream manipulation using common and antithetic random numbers provides a mechanism to induce controlled (and somewhat, predictable) dependency, which in turn can be used to design output performance estimators whose variance is lower than straightforward (crude) estimators based on simple averages.

Control variates can also be designed to induce a known output dependency with the desired performance measures. Control variates can be internal to a simulation model (e.g., an input stream that has a predictable dependency with an output performance measure) or a second simulation model (e.g., a simulation model whose output has a predictable dependency with the output performance measure under study). To ensure that the control variate does not create unwanted bias in the output performance estimator, its expectation should be known. Note that common and/or antithetic random numbers may be used to design the control variate, demonstrating how multiple variance reduction strategies can be simultaneously exploited. An example with single server queues is helpful for students to grasp the design of both internal (e.g., the service time stream with the waiting time in the system output data) and external (e.g., M/M/1 with G/G/1 systems) control variates.

Conditioning by using known information about the model can also lead to lower variance estimators. What conditioning effectively does is allow some of the uncertainty in the simulation output to be reduced based on known information about the system under study. By exploiting such information, one reduces the sample space for the simulation output without sacrificing accuracy or bias. Indeed, the challenge when using conditioning is to identify the appropriate random variable to condition on such that the resulting estimator remains unbiased. Once again, simple queuing system models are helpful for students to visualize how to identify what would be appropriate random variables upon which to condition. An excellent illustrative example for the techniques presented here is described in Nelson (1985).

The key issue when teaching variance reduction techniques is to motivate the value added that such methods bring when analyzing simulation models, and provide practical evidence of their impact on problems.

3 ROBINSON: VERIFICATION AND VALIDATION OF SIMULATION MODELS

3.1 Background

The verification and validation (V&V) of simulation models is obviously a topic of great importance to those developing simulation models and to their clients. The purpose of the teaching session outlined here is to develop an understanding of the concepts of verification and validation, and of the related need to develop confidence in the model so it will be used for decision-making. It also provides a basis for teaching some of the methods that can be used in V&V.

In the session the students are introduced to the key concepts of V&V, they are shown how V&V fits into the modeling life-cycle and we discuss the difficulties that are encountered in verifying and validating simulation models. They are also shown some methods of performing V&V tests on a model. The session finishes with a role-play in which the students act as the clients of a model and I act as the modeler. My job is to convince the students that they should use my model for decision-making. The aim is to actively involve the students in the V&V of a model and to demonstrate some methods of V&V in action. It is on this role-play that we shall largely focus here.

The session is delivered to undergraduate and graduate students on degree programs as part of a detailed module on simulation. It is also delivered to management students (e.g., MBAs) as part of a more general training in the use of models, including simulation. Further, it is included as part of industrial training courses on modeling and simulation. Using the same basic structure, the content of the session can easily be varied to suit the audience, whether they be interested in the technical detail or just the more conceptual level.

3.2 Outline of the Session

The conceptual basis for the session centers on the modeling life-cycle presented in Figure 1 and described in detail in Robinson (2004). This life-cycle is adapted from the work of Landry, Malouin, and Oral (1983), which in turn is based on Sargent's early work on V&V; see Sargent (2011) for his most recent discussion on V&V. Earlier in the course the students will have been introduced to the modeling life-cycle, but without a discussion on V&V (the dashed lines in the figure). In this session we now add V&V to the life-cycle.

The session proceeds as follows:

- Discussion of the definition of the terms verification and validation; a key point is that a model is either valid (it is sufficiently accurate for purpose) or it is not.
- Inclusion of V&V into the simulation life-cycle by building-up Figure 1 and describing the sub-forms of validation denoted in the figure.
- V&V is identified as a continuous process, not a stage, that is performed throughout a simulation study.
- Five key issues in V&V are discussed: there is no such thing as general validity; there is normally no real world to compare the model against; which description of the real world should the model be compared to; where there are real world data for comparison they are often inaccurate or just a sample; there is insufficient time for V&V.
- It is concluded that it is impossible to validate a model, hence the process of V&V is about increasing confidence that a model might be valid, to the point where it will be used for decision-making.
- For a more technical audience a series of techniques for V&V are discussed. The number of methods and technical detail can vary according to the audience and time available.
- The role-play using the Natland Bank case.

Henderson, Jacobson, and Robinson





3.3 The Role-Play: Natland Bank

PowerPoint slides and a version of the simulation model (in the SIMUL8 software) used for this role play are available through www-staff.lboro.ac.uk/~bsslr3. Developing the simulation model in other software is a simple task. For the conceptual content behind this session see Robinson (2004), Chapter 12.

The students are presented with the following problem:

The Natland Bank are opening a new bank branch and want to know how many ATMs to place in the branch. They are currently planning to include two ATMs, but need to know if this will achieve their desired service standard of 95% of customers queueing for less than 3 minutes.

The students are told that they are the bank branch manager and that I have built a simulation model to help them determine the number of ATMs; hence the students are my clients. We then work through the simulation life-cycle together, performing a series of V&V tests on the model that I have developed. The steps involved are now described.

3.3.1 Conceptual Model Validation

I present the modeling objectives, schematic of the model, list the data requirements and highlight the simplifications I have made in the model (e.g., I do not model customers balking). What follows is a discussion with the students (clients) over whether they are happy with the suggested model. This often leads to a range of suggestions for improving the model. It is only after this discussion that I identify that we have been validating the conceptual model by means of communicating the concept to the clients and jointly critiquing the proposed model. One specific point to highlight is that the students probably have quite different levels of confidence in the model, especially as some will accept the model simplifications and others will not.

3.3.2 Data Validation

The students are presented with data on arrival rates and service times. I explain that they, as clients, have provided me with these data. We then discuss the possible sources of the data and the confidence we might have in them. A specific problem we focus on is that given this is a new bank branch there cannot have been any real-world data collection. The option of performing sensitivity analysis with the model is discussed.

3.3.3 Verification/White-Box Validation

I explain that I have verified the model as I built it against the conceptual model. It is now the turn of the students to look at the detail of the model (white-box validation). They now see a computer based running version of the model for the first time. A range of activities can be performed here, but typically we will watch the animation of the model as it is running (face validation), sample some of the model code (e.g., show the students the empirical distributions in the model), and perform an extreme values test (e.g., increase the service time so the traffic intensity is greater than 1).

3.3.4 Black-Box Validation

Since there is no real world to compare this model to (it is a new bank branch), it is not possible to perform a typical black-box validation. However, I demonstrate the concept of comparing the simulation model to another, simpler, model of the problem. This is done in two ways. First the simulation model is simplified by removing all stochastic elements from the model. It is then possible to calculate exactly the performance of the deterministic simulation. For instance, as an alternative model we calculate the teller utilization:

Teller utilization = $\frac{\text{Customers arriving per hour}}{\text{Customers that can be served in an hour}}$

Second, we perform the same calculation above, but with the expected values for customer arrivals and service times derived from the data in the model. Albeit not exact, the simulation should generate output data that are close to the predicted teller utilization values, especially for such a simple model.

The V&V tests carried out in the role-play can be reduced or extended as desired. For instance, the results of the model could be compared to calculations for an M/M/2 queue model; use an exponential distribution in the model for inter-arrival and service times. Queue performance data from another, similar, bank branch could be used for comparative purposes; generate the data using the simulation model with slightly different parameters. This would enable the use of a t-test to be demonstrated. Errors could also be deliberately include in the model and then corrected as part of the role-play.

After each of the above stages I ask the students to reflect on their confidence in the model. At the end, we assess whether the students would use the model for decision-making, i.e. determining how many ATMs to include in the bank. Typically I observe a range of confidence levels, with some students willing to use the model and some not. This is a means for highlighting that confidence is an attribute of the decision-makers, which is derived from their perceptions of the model. We may need to enhance the model to ensure all the decision-makers have sufficient confidence to use the model.

ACKNOWLEDGMENT

We *profusely* thank Barry L. Nelson for the initial idea for this tutorial, for coordinating our efforts in writing and submission, and for producing the final version of this document.

REFERENCES

Cinlar, E. 1975. *Introduction to Stochastic Processes*. Englewood Cliffs, NJ: Prentice-Hall. Cinlar, E. 2011. *Probability and Stochastics*. New York: Springer.

- Landry, M., J. L. Malouin, and M. Oral. 1983. "Model validation in operations research". European Journal of Operational Research 14:207–220.
- Lewis, P. A. W., and G. S. Shedler. 1979. "Simulation of nonhomogeneous Poisson processes by thinning". *Naval Research Logistics* 26 (3): 403–413.
- Nelson, B. L. 1985. "An illustration of the sample space definition of simulation and variance reduction". *Transactions of the Society for Computer Simulation* 2:237–247.
- Resnick, S. I. 1992. Adventures in Stochastic Processes. Boston: Birkhäuser.
- Robinson, S. 2004. Simulation: The Practice of Model Development and Use. Chichester, UK: Wiley.
- Ross, S. M. 1996. Stochastic Processes. 2nd ed. New York: Wiley.
- Rubinstein, R. Y., and B. Melamed. 1998. Modern Simulation and Modeling. Wiley.
- Saltzman, E. A., J. H. Drew, L. M. Leemis, and S. G. Henderson. 2012. "Simulating multivariate nonhomogeneous Poisson processes using projections". *ACM Transactions on Modeling and Computer Simulation* To appear.
- Sargent, R. G. 2011, December. "Verification and validation of simulation models". In *Proceedings of the* 2011 Winter Simulation Conference, edited by S. Jain, R. R. Creasey, J. Himmelspach, K. P. White, and M. Fu, 183–198. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

AUTHOR BIOGRAPHIES

SHANE G. HENDERSON is a professor in the School of Operations Research and Information Engineering at Cornell University. His research interests include discrete-event simulation and simulation optimization, and he has worked for some time with emergency services. He is the current chair of the INFORMS Applied Probability Society, and an associate editor for both Management Science and Stochastic Systems. He co-edited the Proceedings of the 2007 Winter Simulation Conference. His web page can be found at people.orie.cornell.edu/~shane.

SHELDON H. JACOBSON is a Professor and Director of the Simulation and Optimization Laboratory at the University of Illinois at Urbana-Champaign. He has a B.Sc. and M.Sc. (both in Mathematics) from McGill University, and a M.S. and Ph.D. (both in Operations Research and Information Engineering) from Cornell University. He has been a member of I-SIM (and its predecessor, the College on Simulation) since 1988, serving as its Treasurer from 1994 to 1996, on the Best Publication Award Committee from 1997 to 1999 (chairing the committee in 1998), and the Membership Committee from 2000 to 2002. He also served as an Associate Editor for Operations Research, Simulation Department from 2001-2005. He presently serves in a number of editorial roles, including the Focused Issue Editor for IIE Transactions, Operations Engineering and Analysis (since 2009). A regular attendee of the Winter Simulation Conference and INFORMS national meetings, his research has been disseminated in a wide spectrum of journals, as well as in two research videos and through the popular media. He is a member of INFORMS, IIE, ASEE, and SIAM.

STEWART ROBINSON is Professor of Management Science and Associate Dean for Research at Loughborough University, School of Business and Economics. Previously employed in simulation consultancy, he supported the use of simulation in companies throughout Europe and the rest of the world. He is author/co-author of five books on simulation. His research focuses on the practice of simulation model development and use. Key areas of interest are conceptual modeling, model validation, output analysis and alternative simulation methods (discrete-event, system dynamics and agent based). He has recently completed a research project that investigated the use of simulation with lean in healthcare. Professor Robinson is co-founder of the Journal of Simulation and Vice President of the Operational Research Society. Home page: www.btinternet.com/~stewart.robinson1/sr.htm