Proceedings of the 1999 Winter Simulation Conference
P. A. Farrington, H. B. Nembhard, D. T. Sturrock, and G. W. Evans, eds.

SMG: A NEW SIMULATION/OPTIMIZATION APPROACH FOR LARGE-SCALE PROBLEMS

Christopher W. Zobel

Department of Management Science
and Information Technology
Virginia Polytechnic Institute and State University
Blacksburg, VA 24061, U.S.A.

William T. Scherer

Department of Systems Engineering
University of Virginia
Charlottesville, VA 22903, U.S.A.

ABSTRACT

It typically can be difficult to create and solve optimization models for large-scale sequential decision problems, examples of which include applications such as communications networks, inventory problems, and portfolio selection problems. Monte Carlo simulation modeling allows for the creation and evaluation of these large-scale models without requiring a complete analytical specification. Unfortunately, optimization of such simulation models is especially difficult given the large state spaces that often produce a combinatorially explosive number of potential solution policies.

In this paper we introduce a new technique, Simulation for Model Generation (SMG), that begins with a simulation model of the system of interest and then automatically builds and solves an underlying stochastic sequential decision model of the system. Since construction and implementation of the created model requires approximation techniques, we also discuss several types of error that are induced into the decision process. Fortunately, the decision policies produced by the SMG approach can be directly evaluated in the original simulation model - thus the results of the SMG model can be compared against any other possible strategies, including any decision policies currently in use.

1 INTRODUCTION

In this paper, we present a new simulation/optimization approach to the generation of optimal policies for very large-scale systems. The approach, which we call the Simulation for Model Generation (SMG) algorithm, combines Monte Carlo simulation with state space aggregation to empirically create an aggregated sequential stochastic decision model (Markov decision process - MDP) representation of the system of interest. This aggregated MDP can then be solved to produce a policy solution for the original system. The basic SMG algorithm is described below, followed by a brief discussion of some of the modeling and simulation issues involved in its implementation.

2 SMG ALGORITHM

One way in which to deal with problems with very large, complex state spaces is to formulate a simpler aggregated model which can then be solved to generate a solution for the original system (see Scherer and White (1986), for an example). This formulation, however, requires combining the parameters from the original model in some way to create the transition probabilities and one-step rewards for its aggregated counterpart (see Puterman (1994)). If the original state and action spaces are very large, for example on the order of 1e+100 elements, then it is not feasible to use an analytic approach to form this combination. The SMG approach allows us to deal with such a situation in a natural way.

Assuming that there exists a verified and validated simulation model of the system under consideration, the SMG approach involves specifying a state space aggregation scheme and identifying the associated mapping from each original state into its aggregated counterpart. By then simulating the underlying process, we may empirically capture the probability of each transition between <u>aggregated</u> states, as well as the associated transition reward or cost. After sufficient "training" of the aggregated parameters, we may solve the associated aggregated MDP model to generate a solution policy which may be "disaggregated" and applied to the original non-aggregated system.

The relative performance of a policy generated by this SMG approach is easily validated by applying it within the "ground-truth" simulation of the system. This ensures that the performance of any solution policy can be accurately determined regardless of how accurate and/or appropriate the chosen aggregation scheme may be. Also, the issue of a poorly chosen aggregation scheme may be addressed through the "Outer loop" of the algorithm, as presented in Figure 1. This provides a framework within which an

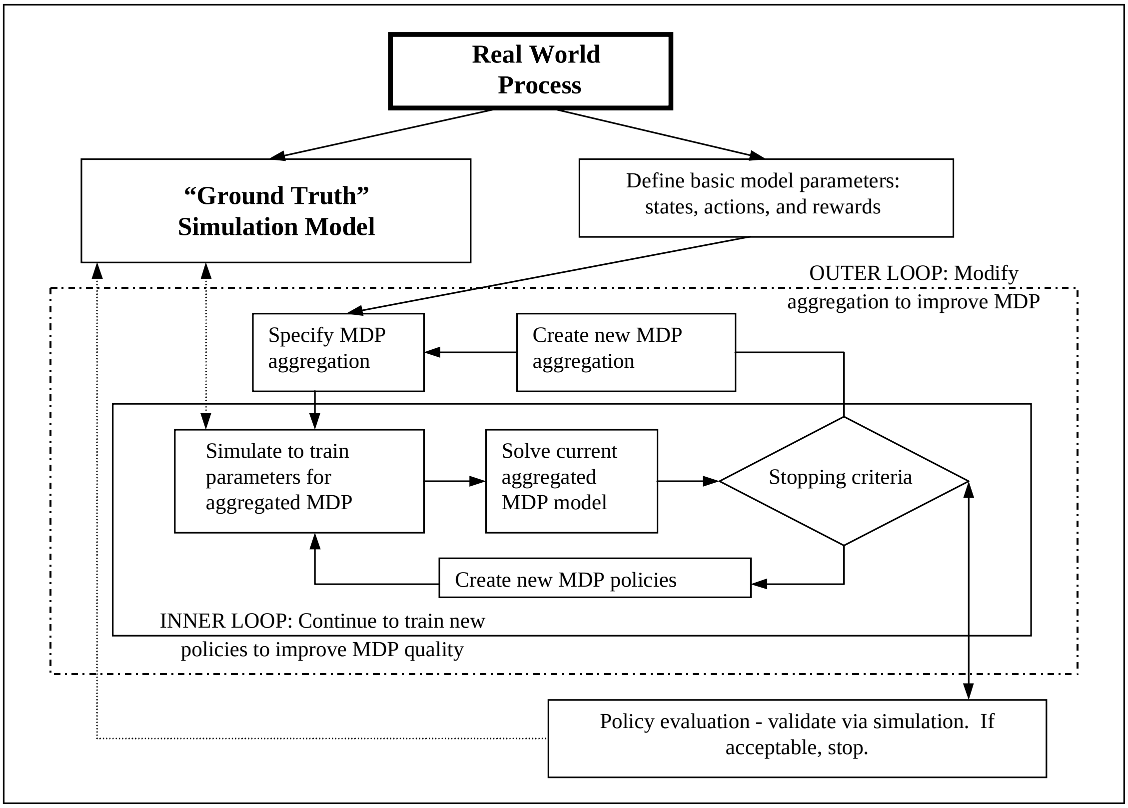

Figure 1: Simulation for Model Generation (SMG)

aggregation scheme may be changed and then reapplied within the basic SMG algorithm to produce a new, and possibly improved, solution policy for the original problem.

3 MODELING AND SIMULATION ISSUES

There are several types of approximation error involved with implementing an approach such as the SMG algorithm, and it is important to recognize these potential errors so that they may be addressed appropriately. The following are the three primary sources of error associated with the SMG algorithm in particular:

1. Type A error - associated with the initial modeling of the underlying problem as a Markov decision process,
2. Type B error - associated with the structural limitations imposed by the formation of an aggregated model and the potential loss of the Markov property due to this aggregation,
3. Type C error - associated with the incomplete simulation training of the generated MDP model parameters.

Each of these error types is illustrated in Figure 2 on the following page.

The first type of error is related to the initial formation of the full-scale Markov decision process from the original underlying problem. Although we typically assume that this original problem has a natural representation as an MDP, this need not be true in general. By carefully defining state and action spaces and by making assumptions about the associated transition behavior, it is possible to model a great number of problems as Markov decision processes, even if they are non-Markovian.

The second type of application error, Type B, is related to the formation of an aggregated model from the original Markov decision process model. There are two aspects to this type of error, both of which impact the ability of the aggregation scheme to accurately represent the underlying process. The first of these is related to the structural limitations that may be imposed upon an aggregated problem due to combining the states from the original model. The second aspect of the Type B error is associated with the potential loss of the Markov property due to the choice of aggregation scheme, and the uncertainty that this imposes on the performance of the aggregated model (Zobel 1998).

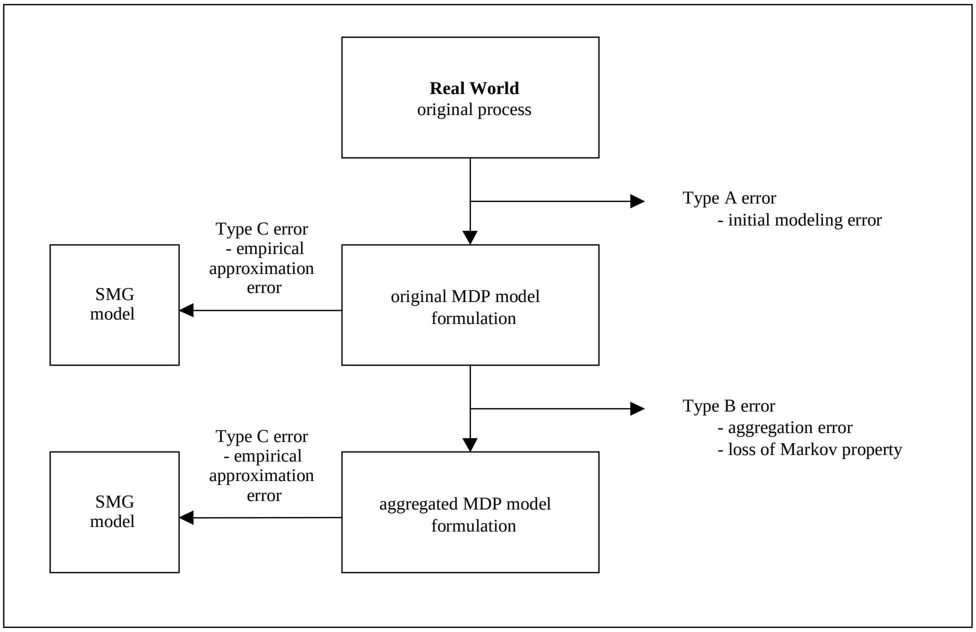

Figure 2: Approximation Error for the SMG Algorithm

The Type C approximation error introduced above is associated with our use of simulation to empirically generate an MDP model. The empirical parameter estimates produced in each iteration of the SMG algorithm are simply approximations to the "true" parameter values for the chosen aggregation, therefore, although these approximations become more accurate as the number of simulated observations increases, they may only achieve the actual values in the limit. By solving a Markov decision process based on these approximate parameters, we are thus subject to the error induced by this result.

In general, it is important to keep in mind that if an SMG generated model is based on an aggregated MDP formulation then it may contain components of all three types of approximation error discussed above. In defining each aspect of the SMG approach, therefore, we must be aware of the presence and the effect of this error so that we may minimize its impact. Ultimately, however, we judge the performance of the algorithm based upon the performance of the policies that it produces within the ground truth simulation model.

4 ILLUSTRATION

We have implemented the SMG approach within two particular test domains: the first is a large-scale telephone network routing problem, and the second is a common, but theoretically and computationally complex, problem of inventory control. We briefly describe each of these problems below.

4.1 Telephone Network Routing Problem

We first explored the performance of the SMG algorithm with respect to the problem of choosing routes for telephone calls within a network. Each call is associated with an origin node and a destination node and occupies a single connecting path through the network for the (stochastic) duration of the call. Furthermore, each link between nodes only has finite capacity and the arrival of calls to the network occurs stochastically. The state of the system is defined to be the number and location of calls in the network, an action is a routing choice for a particular incoming call, and the objective is to minimize the long-term number of calls "blocked" or prevented from entering the network because of insufficient capacity on the specified route.

Using several different network formulations, the results of applying the SMG algorithm to the routing problem were compared against four standard heuristic routing strategies: direct routing, least-loaded routing, dynamic alternative routing, and Krishnan's MDP based Separable Routing (1990). Despite the complexity of the problem, the SMG approach was able to generate telephone routing policy solutions that compared favorably against these well-established, tailored heuristics. In several situations, SMG actually matched both the performance and the structure of the best heuristic solution, without previous knowledge of this structure. Additional testing will continue to improve the performance of the algorithm and will explore some of the issues associated with aggregating the network state space.

4.2 Inventory Control Problem

We chose as our second application a standard inventory control problem, with discrete stochastic monthly demands and variable storage, shortage, and ordering costs. In this example, the current inventory level defines the state of the system, the actions each correspond to a given number of items to order in a given month, and the objective function involves minimizing the average expected cost per unit time. Aggregated problems are easily generated from this model by combining adjacent inventory levels.

As the basis for our SMG approach, we used a "black-box" simulation of this inventory control problem which was drawn from the literature (Law and Kelton 1991), and which provided the algorithm only with the current system state and transition cost at each decision interval. With no other system information available, the SMG algorithm was able to generate policy solutions that were of the complex structure and quality of the known optimal solution form, i.e., the well-studied (s, S) policy structure (Zobel 1998). This suggests that the approach may be very useful in situations where there is limited system knowledge and where the optimal policies take on a specific, but unknown, structure.

5 CONCLUSIONS

In this paper we have presented an approach by which we may take a problem that is computationally intractable and create reasonable solutions via a simulation of the original problem. The solution methodology described uses a simulation model to create an approximating MDP which is then solved via traditional MDP solution approaches. Since the approximating MDP model is a fairly rich and robust sequential optimization model, optimal policies can be created which represent an intelligent and comprehensive search of the policy space.

The performance of policies generated by the SMG algorithm may be properly evaluated by applying them within an underlying “ground-truth” simulation which has itself been verified and validated. In this way, any solution generated by SMG can be easily compared against any other potential solution via the simulation model. Exploratory experimental results indicate that the SMG algorithm would allow for the modeling and optimization of numerous problems for which no known policy structure exists, and which are currently considered too complex for anything other than simple heuristic rules. It is our belief that these results demonstrate that this approach is of high potential and worthy of additional research and testing.

REFERENCES

Krishnan, K. R. 1990. Markov Decision Algorithms for Dynamic Routing, *IEEE Communications Magazine*, 28:10, 66-69.

Law, A. L., and W. D. Kelton. 1991. *Simulation Modeling and Analysis*. 2d ed. New York: McGraw-Hill.

Puterman, M. L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. New York: John Wiley & Sons.

Scherer, W. T., and C. C. White, III. 1986. A Planning and Decision-Aiding Procedure for Purchasing and Launching Spacecraft. *Interfaces* 16:3, 31-40.

Zobel, C. 1998. An Investigation into the Use of Simulation for Solving Large-scale Markov Decision Processes. Ph.D. dissertation, Department of Systems Engineering, University of Virginia, Charlottesville, Virginia.

AUTHOR BIOGRAPHIES

CHRISTOPHER W. ZOBEL is an Assistant Professor of Management Science and Information Technology at the Virginia Polytechnic Institute and State University. He holds a Ph.D. in Systems Engineering from the University of Virginia, an M.S. in Mathematics from the University of North Carolina at Chapel Hill, and a B.A. in Mathematics from Colgate University. Professor Zobel’s primary research interests are in the areas of large-scale stochastic decision problems, intelligent decision systems, heuristic problem-solving, and simulation. He is a member of INFORMS and DSI.

WILLIAM T. SCHERER is an Associate Professor of Systems Engineering at the University of Virginia, where he also received the B.S., M.E., and Ph.D. degrees in Systems Engineering. Professor Scherer's research interests include scheduling systems, decision analysis, intelligent systems, combinatorial optimization, Markov decision processes, and strategic planning and intelligent transportation systems. He is a member of IEEE, INFORMS, and DSI.