Proceedings of the 2023 Winter Simulation Conference
C. G. Corlu, S. R. Hunter, H. Lam, B. S. Onggo, J. Shortle, and B. Biller, eds.

CAUSAL DYNAMIC BAYESIAN NETWORKS FOR SIMULATION METAMODELING

Pracheta Amaranath

Manning College of Information and Computer Sciences
University of Massachusetts Amherst
140 Governors Drive
Amherst, MA 01003, USA

ABSTRACT

A traditional metamodel for a discrete-event simulation approximates a real-valued performance measure as a function of the input-parameter values. We introduce a novel class of metamodels based on *modular dynamic Bayesian networks* (MDBNs), a subclass of probabilistic graphical models which can be used to efficiently answer a rich class of *probabilistic and causal queries* (PCQs). Such queries represent the joint probability distribution of the system state at multiple time points, given observations of, and interventions on, other state variables and input parameters. This paper is a first demonstration of how the extensive theory and technology of causal graphical models can be used to enhance simulation metamodeling. We demonstrate this potential by showing how a single MDBN for an M/M/1 queue can be learned from simulation data and then be used to quickly and accurately answer a variety of PCQs, most of which are out-of-scope for existing metamodels.

1 INTRODUCTION

Simulation metamodels approximate the behavior of a simulation model, avoiding the need to run computationally expensive experiments. Discrete-event simulation metamodels have traditionally focused on computing statistical relationships between a set of input parameters and a single output measure of interest. For example, a simulation metamodel of an M/M/1 queue might approximate the expected time-average queue length over a time interval, or the probability of the queue length at a given time, as a function of the arrival and service rates of the queue. Unfortunately, separate metamodels must be constructed for each query of interest, requiring substantial computational effort. Moreover, separate metamodels for similar queries on the same system may be mutually inconsistent.

This paper presents a powerful, complementary approach to constructing simulation metamodels by applying tools and techniques developed by researchers in graphical models and causal inference. Specifically, we show how *modular dynamic Bayesian networks* (MDBNs), a subclass of probabilistic graphical models (Koller and Friedman 2009), can be used to efficiently estimate answers to a rich class of *probabilistic and causal queries* (PCQs). PCQs are conditional probability distributions of the variables in a simulation model. They represent the joint probability distribution of the system state at multiple time points, given observations of, and interventions on, system states and simulation parameters at other time points. PCQs can also include “inverse” queries, which treat simulation parameters as Bayesian random variables and represent the probability distribution of the parameters given observations of system states. The results from such queries can be used, for example, to select the input parameters that would maximize the probability of achieving target system-state values at specific times.

An MDBN is a dynamic Bayesian network (Murphy 2002) that is intended to approximate the evolution of the simulation states over multiple times of interest. A *single*, trained MDBN can answer a broad range of PCQs, thereby enabling useful, efficient, and fine-grained interrogation of the dynamic behavior of a simulation model that cannot be achieved via traditional simulation metamodels alone. Importantly, the constructed MDBN need not be trained by explicitly simulating interventions on the simulation model. Instead, interventional queries can be quickly answered via exact or approximate inference methods, which are well-studied techniques in graphical models.

2 METHODOLOGY

To create and use an MDBN metamodel of an M/M/1 queue, we first specify a structure for the MDBN that includes all simulation states of interest. Figure 1a is a depiction of an MDBN for the M/M/1 queue for M discrete time steps with λ and μ representing the arrival and service rates respectively. We use $L^{(j)}$ to represent the queue length at time instant j. Next, we estimate the conditional probability distributions (CPDs) of the MDBN using data collected by running experiments on the simulation model. For the above metamodel, we need to estimate (or specify fixed distributions) for the following CPDs: $\mathbb{P}_1 = P(\lambda), \mathbb{P}_2 = P(\mu), \mathbb{P}_3 = P(L^{(0)})$ and $\mathbb{P}_4 = P(L^{(j+1)}|\lambda, \mu, L^{(j)})$. This MDBN can now be queried to estimate answers to PCQs of interest. In brief, we do this by translating a PCQ into an equivalent intervention-free form, if needed, and then applying standard inference algorithms for probabilistic BNs.

3 RESULTS

We demonstrate three classes of PCQs that can be answered using the MDBN: (1) *extrapolative queries* that infer the probability distribution of queue length for time $\tau_i > T$, where T is the time horizon of the simulations in the training data, (2) *interventional queries* that infer probability distributions of a specified set of simulation variables given modifications to the values of a different set of simulation variables, and (3) *inverse queries*. We find that the MDBN accurately captures the distribution of the queue length L when compared to running multiple, computationally intensive experiments using the simulator. Figure 1b depicts the Jensen-Shannon distance between the probability distribution of $P(L_5)$ upon an intervention to the queue length at L_4 computed via the MDBN and the simulator. Figure 1c is a demonstration of the results of an inverse query, where we find the value of λ that maximizes the probability of being in a target set of states at a specified time, without having to evaluate each possibility by running simulations. Even though there are some disparities between the exact simulation and the MDBN metamodel, the metamodel selects the optimal value of λ without the need for expensive simulation runs.

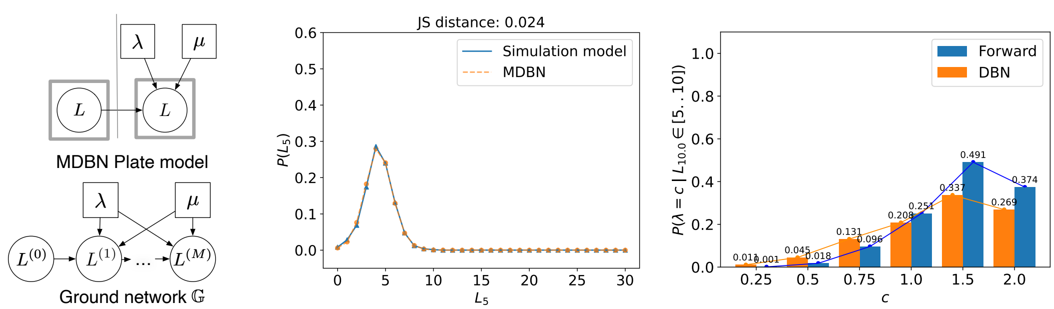

(a) MDBN: M/M/1 queue (b) Intervention: $L_4 \leftarrow L_4 + 4$. (c) Optimizing λ to control $L_{10.0}$.

Our work provides an entry point to the highly expressive family of probabilistic graphical models (PGMs). This family of models has the potential to greatly expand the scope of simulation metamodeling.

REFERENCES

Koller, D., and N. Friedman. 2009. *Probabilistic Graphical Models: Principles and Techniques*. Cambridge, Massachusetts: MIT Press.

Murphy, K. P. 2002. *Dynamic Bayesian Networks: Representation, Inference and Learning*. Ph.D. thesis, University of California, Berkeley, California.