Recent Advances in Importance Sampling for Statistical Model Checking

Daniël Reijsbergen Pieter-Tjerk de Boer Werner Scheinhardt Boudewijn Haverkort
University of Twente, Enschede, The Netherlands
{d.p.reijsbergen,p.t.deboer,w.r.w.scheinhardt,b.r.h.m.haverkort}@utwente.nl

— Extended abstract for SMC 2013 —

In the following work we present an overview of recent advances in rare event simulation for model checking made at the University of Twente. The overview is divided into the several model classes for which we propose algorithms, namely multicomponent systems, Markov chains and stochastic Petri nets, and probabilistic timed automata.

1 Introduction

Probabilistic model checking is an increasingly popular technique for the analysis of the performance of highly reliable systems. Using probabilistic model checking, one exhaustively analyses the behaviour of a system until a statement can be made about the probability that, or the time until some interesting event occurs in the system. If the state space of the model is very large, which is often the case for realistic systems, exact numerical evaluation techniques are computationally infeasible. In such cases, statistical model checking techniques based on discrete-event simulation of the system models are a viable alternative. However, although in principle always applicable, discrete-event simulation comes with a number of challenges, *rare events* being one of them: if the probability that the event of interest occurs is very low, the number of runs needed to even witness the event can be impractically large. In highly reliable systems one is often interested in estimating small failure probabilities, meaning that efficient simulation techniques meant for rare events for model checking are a vital addition to the toolset of the model checker.

At the University of Twente we combine our solid knowledge in both the field of probabilistic and stochastic model checking and the field of rare event simulation. Research on this cross-over subject has recently led to new results for fully Markovian models, as well as to new results for models described as probabilistic timed automata (PTA). We briefly overview these results in this short paper.

2 Multicomponent Systems

In [7], we investigate rare event simulation for highly reliable Markovian multicomponent systems, a well-known case study from the literature. These systems consist of several independent component types such that the components of each type act as spares for the others of the same type. The system as a whole breaks down if several components of a single type break down. We are interested in the probability of observing a system failure before a specific time bound, and we assume that this probability is very low. Our research shows that there are two interesting parametric regimes that cause system failure to be a rare event: component failures can be slow or component repairs can be fast. In the former

Submitted to:
SMC 2013

setting, an importance sampling scheme similar to the one of [2] works well. This scheme is based on the *dominant paths to failure*; we assume that the total probability of failure is close to the probability of a series of subsequent component failures. Since this is a valid assumption in the former setting — after all, if failures are slow then it is unlikely to witness several of them before the time bound — the importance sampling scheme works well. However, in the latter regime it does not, because several component failures almost immediately followed by repairs are in this setting not unlikely.

As a remedy for the fast-repair scenario, we apply a renewal argument, that is, we partition an execution of the system model into disjoint *busy cycles*. The number of busy cycles needed before we reach system failure has a *geometric* distribution, and the duration of each busy cycle is approximately exponentially distributed. From elementary probability theory we then know that the time until system failure is also approximately exponentially distributed; we use this knowledge to construct a *state*- and *time*-dependent importance sampling scheme. Under this scheme, it may still be likely that the first busy cycle ends without system failure, but as time passes the probability of system failure during the current busy cycle increases. We empirically demonstrate that this importance sampling scheme works well.

In subsequent research we tackled the problems that arise when the state in which all components are up is part of a so-called *high-probability cycle* (more on that in the next section). In this setting, it can be computationally expensive to use the same notion of busy cycles as previously. We showed in [7] that in the relevant parametric regime the time until system failure can still be seen as exponentially distributed, meaning that our technique was still applicable.

3 General Markov chains and Petri Nets

In [6] and [8] we generalised the dominant paths approach to general Markov chains and stochastic Petri nets. The idea is to construct an algorithm that derives a well-performing importance sampling scheme in an automated fashion, while requiring as little information about the behaviour of the system as possible from the investigator. Typically, in these models we assume that each transition has a probability that behaves like $\varepsilon > 0$ to some power (possibly 0), where ε is the rarity parameter which is common to all transitions. Transitions for which the power is zero are 'high-probability' transitions (typically corresponding to repairs), while all others become rare as ε becomes small, possibly at different rates (typically corresponding to different types of failures). The event of interest in this setting is the event that we reach some rare set of states before another typical set. The dominant paths to failure are the paths that lead to the rare set and which have the lowest power of ε. The idea is that if we can find the dominant paths, we can use the results from [3] to construct an importance sampling scheme which has the desirable property of *bounded (or even vanishing) relative error*, which means that if ε is made smaller, the number of samples needed to achieve some relative level of accuracy remains bounded (or even decreases). The key to finding the dominant paths is to find a *distance function* that assigns to each state in the model the relative distance of that state to the rare failure state, in terms of powers of ε.

In [8] we propose a method for automatically finding the dominant paths to failure that works on the state space of the Markov chain. We find these distances using an application of Dijkstra's algorithm. The dominant path probabilities can then be found by back-tracking the algorithm. In the meantime, we are able to remove high-probability cycles, which are a theoretical hurdle to apply the results of [3]. A high-probability cycle occurs if a state exists that can reach itself by a path of non-zero length that consists of only high-probability steps. We use lumping techniques known from the model checking community to remove these cycles by grouping them into a single state.

In [6] we take our methodology one step further and work on the level of the stochastic Petri net.

Using this approach, the state space can be divided into *zones* such that for all states in a zone, the dominant paths to a set of rare failure states have the same form. The zones are demarcated by a set of affine inequalities, thus avoiding the need to enumerate all states individually. The number of zones in typical models does not need to increase as the number of states in the model increases.

4 Probabilistic Timed Automata

In a recent study, we investigated how typical rare-event scenarios and importance sampling recipes as developed for full stochastic (Markovian) models, can apply also in the context of statistical model checking for PTA.

First of all, we investigated different types of rare events, as they can occur in PTA. The most important ones are due to one of the three following constructs in a PTA: (i) the PTA features probabilistic branches with largely different (competing) probabilities, e.g., a probability 10^{-5} to go to a failure state, and $1-10^{-5}$ to go to a benign state; (ii) the PTA features a "series" of locations, each with a large numbers of equiprobable outgoing edges, leading to a very small overall probability to reach certain locations; (iii) the timing (distribution) in locations is such that the probability to reach certain locations "behind" these, becomes very small. Other, more intricate rare events can be constructed, certainly if one makes full use of PTA-language, such as provided in Uppaal [1], however, we consider the above three as the most important ones. Each of these three types can cause difficulties in estimating the time until a certain event takes place, or in estimating the probability that a certain event takes place (with or without time bound).

The usefulness of the importance sampling schemes failure biasing [9] and delay biasing (which in turn is similar to forcing [4]), has subsequently been investigated experimentally, on a limited number of small models. For **failure biasing**, in the context of Markovian models, two variants have been proposed: simple failure biasing and balanced failure biasing (BFB), where the latter appears most suitable to be used for handling PTA with largely differing probabilities on outgoing edges. A new BFB scheme for PTA is proposed, inspired by the Markovian BFB scheme, that adapts failure and repair probabilities (when these notions can be sensibly attributed to the edges), as well as the probabilities for selecting loop edges. The corresponding likelihood ratio can easily be computed. A limited number of case studies is performed, and shows reasonable results for variance reduction, that is, sample standard deviations being a factor 2 up to 1000 smaller. However, also cases were seen where, in fact, a variance increase occurred. Hence, although some good results could be obtained with BFB, more work is needed to come to more reliable changes of measure.

With **delay biasing**, one changes the state/location residence time distribution such that paths to failure events (locations) are more often selected. First point of care in this case is to deal with so-called guard gaps and guard ranges, that is, we have to be precise about which time delays values are actually allowed given the composition of the location invariant and the ranges on the outgoing edges. Taking care of these, we have developed two new delay biasing schemes, one in case the location supports a uniform delay, and one in case the delay in a location follows an exponential distribution. The corresponding likelihood ratios can easily be expressed. The results on a set of 5 small case studies show a reduction of estimator standard deviation of up to one order of magnitude.

A simple tool has been develop, that allows us to simulate, using our new importance sampling techniques, simple PTA models. These models can be described using Uppaal; using comments (from the viewpoint of Uppaal), "instructions" can be given to our tool, that uses the XML-file, as generated by Uppaal, as its input.

5 Acknowledgements

This work has been supported by the Netherlands Organisation for Scientific Research (NWO) in their project 612.064.812. We thank Marco Pasveer for his contributions, as laid down in his M.Sc. thesis [5], and dr. M. Stoelinga for her contribution to [5].

References

[1] J. Bengtsson, K. Larsen, F. Larsson, P. Pettersson, and W. Yi. UPPAAL — a tool suite for automatic verification of real-time systems. *Hybrid Systems III*, pages 232–243, 1996.

[2] P.T. de Boer, P. L'Ecuyer, G. Rubino, and B. Tuffin. Estimating the probability of a rare event over a finite time horizon. In *Proceedings of the 2007 Winter Simulation Conference*, pages 403–411, 2007.

[3] P. L'Ecuyer and B. Tuffin. Approximating zero-variance importance sampling in a reliability setting. *Annals of Operations Research*, 189(1):277–297, 2011.

[4] V. Nicola, P. Shahabuddin, and M. Nakayama. Techniques for fast simulation of models of highly dependable systems. *IEEE Transactions on Reliability*, 50(3):246–264, 2001.

[5] M. Pasveer. Importance sampling for probabilistic timed automata. Master's thesis, University of Twente, 2013.

[6] D. Reijsbergen, P.T. de Boer, W. Scheinhardt, and B. Haverkort. Automated rare event simulation for stochastic Petri nets. Accepted for publication at QEST 2013.

[7] D. Reijsbergen, P.T. de Boer, W. Scheinhardt, and B. Haverkort. Rare event simulation for highly dependable systems with fast repairs. *Performance Evaluation*, 69(7):336–355, 2012.

[8] D. Reijsbergen, P.T. de Boer, W. Scheinhardt, and S. Juneja. Some advances in importance sampling of reliability models based on zero variance approximation. In *Proceedings of RESIM 2012*. NTNU University Press.

[9] P. Shahabuddin. Importance sampling for the simulation of highly reliable Markovian systems. *Management Science*, pages 333–352, 1994.