Planning and Verification for Stochastic Processes with Asynchronous Events

Håkan L. S. Younes
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213, USA
lorens@cs.cmu.edu

Abstract

We consider a general model of stochastic discrete event systems with asynchronous events, and propose to develop efficient algorithms for verification and control of such systems.

Introduction

Dynamic systems consisting of concurrent interacting processes present a broad spectrum of uncertainty, including uncertainty in the occurrence of exogenous events and the outcome of actions selected by a control policy. In this work, we are concerned with stochastic discrete event systems with asynchronous events, and where the delay from when an event is enabled until it triggers is governed by an arbitrary positive distribution.

As an example, consider a network of computers. Each computer can be in one of two states: working or failed. Failure is caused by the triggering of an exogenous event, with the time from when a computer starts working until it fails being a random variable associated with some positive distribution, for example a Weibull distribution.

In addition to exogenous events, we can have controllable events (actions) that can be enabled or disabled at will by a control policy. For the network of computers, we could introduce a reboot action for each computer that can restore a failed computer to the working state. Given a system of this type, we are interested in generating a control policy for the system satisfying some objective, which can be to either satisfy a formula in a temporal stochastic logic or to maximize expected reward.

We conjecture that asynchronous stochastic decision processes with general delay distributions exhibit sufficient structure as to allow computationally tractable algorithms for policy generation. We will prove this by developing efficient algorithms for probabilistic and decision theoretic planning with complex stochastic domain models. As part of this effort, we will also consider the problem of verifying properties of stochastic discrete event systems, and we will develop scalable verification algorithms for such systems.

The problem of planning under uncertainty has been addressed by researchers in operations research, artificial intelligence (AI) and control theory, among others. Numerous approaches have been proposed, but current methods for planning under uncertainty cannot handle the full generality of the model we are considering. Guestrin, Koller, & Parr (2002) have shown that factored Markov decision processes (MDPs) can be used to handle concurrent actions, but this approach is restricted to sets of actions executed in synchrony. Continuous-time MDPs (Howard 1960) can be used to model asynchronous systems, but are restricted to events and actions with exponential delay distributions. Semi-Markov decision processes (SMDPs) (Howard 1971) lift the restriction on delay distributions, but cannot model asynchrony.

We adopt the generalized semi-Markov process (GSMP), first introduced by Matthes (1962), as our domain model. The GSMP is an established formalism in queuing theory for modeling continuous-time stochastic discrete event systems (Glynn 1989). Intuitively, a GSMP can be viewed as the composition of concurrent semi-Markov processes, and is a natural model for systems with asynchronous components.

Accomplishments to Date

The idea of probabilistic planning with GSMP domain models was first explored by Younes & Musliner (2002) in an effort to extend CIRCA (Musliner, Durfee, & Shin 1995)—an architecture for real-time control—to handle stochastic information in the form of probability distributions over event trigger times. CIRCA generates control policies incrementally, using a verifier module to ensure that each expansion of a policy remains safe relative some safety constraints specified by the user. Younes & Musliner developed a plan verification algorithm for GSMP domain models based on sequential acceptance sampling (Wald 1945).

The plan verification algorithm was later generalized by Younes & Simmons (2002) to handle time-bounded properties expressed in CSL (Aziz *et al.* 2000; Baier *et al.* 2003), a logic formalism for expressing temporal and probabilistic properties of stochastic systems. Younes *et al.* (2004) compare this statistical approach for CSL model checking with numerical solution techniques for continuous-time Markov chains based on uniformization (Baier *et al.* 2003) and show that the statistical approach scales better with an increase in the size of the state space, while the numerical approach is less sensitive to decreases in error tolerance.

A general framework for policy generation in continuous-time stochastic domains with asynchronous events and actions was presented by Younes, Musliner, & Simmons (2003) based on the Generate, Test and Debug paradigm (Simmons 1988). The test step is carried out using the statistical verification algorithm mentioned above making it possible to plan for arbitrary time-bounded CSL goals. Concrete techniques for representing and generating initial policies using an existing deterministic temporal planner, robust sample path analysis techniques for extracting failure scenarios, and efficient policy repair techniques were presented in a followup paper (Younes & Simmons 2004a). These techniques have been implemented in TEMPASTIC, a novel planner accepting domain descriptions in an extension of PDDL developed by Younes (2003) and using VHPOP (Younes & Simmons 2003) as a deterministic planner.

Remaining Research Effort

While TEMPASTIC represents a step towards practical planning algorithms for complex stochastic domains with asynchronous events, there is still work to be done to improve its efficiency. We need to develop heuristics for selecting failure scenarios that best will guide the search towards a policy satisfying a given CSL goal. We are also exploring ways to improve the efficiency of the verification algorithm, including distributed algorithms.

In addition to research on probabilistic planning, we are pursuing decision theoretic approaches as well. For this purpose, we will introduce the generalized semi-Markov decision process (GSMDP) as a model for decision theoretic planning with asynchronous events and actions. We need to develop practical policy generation algorithms for GSMDPs, and are currently exploring various possible directions for this research. In a recent paper (Younes & Simmons 2004b) we propose using phase-type distributions (Neuts 1981) to approximate general distributions in the GSMDP model, which enables us to use standard MDP techniques to approximately solve GSMDPs. We also expect Monte Carlo integration and value function approximation using nearest neighbor learning, previously used for solving partially observable MDPs (Thrun 2000), to be useful for solving GSMDPs.

References

Aziz, A.; Sanwal, K.; Singhal, V.; and Brayton, R. 2000. Model-checking continuous-time Markov chains. *ACM Transactions on Computational Logic* 1(1):162–170.

Baier, C.; Haverkort, B. R.; Hermanns, H.; and Katoen, J.-P. 2003. Model-checking algorithms for continuous-time Markov chains. *IEEE Transactions on Software Engineering* 29(6):524–541.

Glynn, P. W. 1989. A GSMP formalism for discrete event systems. *Proceedings of the IEEE* 77(1):14–23.

Guestrin, C.; Koller, D.; and Parr, R. 2002. Multiagent planning with factored MDPs. In *Advances in Neural Information Processing Systems 14: Proc. 2001 Conference*. Cambridge, MA: The MIT Press. 1523–1530.

Howard, R. A. 1960. *Dynamic Programming and Markov Processes*. New York, NY: John Wiley & Sons.

Howard, R. A. 1971. *Dynamic Probabilistic Systems*, volume II. New York, NY: John Wiley & Sons.

Matthes, K. 1962. Zur Theorie der Bedienungsprozesse. In *Trans. Third Prague Conference on Information Theory, Statistical Decision Functions, Random Processes*, 513–528. Publishing House of the Czechoslovak Academy of Sciences.

Musliner, D. J.; Durfee, E. H.; and Shin, K. G. 1995. World modeling for the dynamic construction of real-time control plans. *Artificial Intelligence* 74(1):83–127.

Neuts, M. F. 1981. *Matrix-Geometric Solutions in Stochastic Models: An Algorithmic Approach*. Baltimore, MD: Johns Hopkins University Press.

Simmons, R. G. 1988. A theory of debugging plans and interpretations. In *Proc. Seventh National Conference on Artificial Intelligence*, 94–99. AAAI Press.

Thrun, S. 2000. Monte Carlo POMDPs. In *Advances in Neural Information Processing Systems 12: Proc. 1999 Conference*. Cambridge, MA: The MIT Press. 1064–1070.

Wald, A. 1945. Sequential tests of statistical hypotheses. *Annals of Mathematical Statistics* 16(2):117–186.

Younes, H. L. S., and Musliner, D. J. 2002. Probabilistic plan verification through acceptance sampling. In *Proc. AIPS-02 Workshop on Planning via Model Checking*, 81–88.

Younes, H. L. S., and Simmons, R. G. 2002. Probabilistic verification of discrete event systems using acceptance sampling. In *Proc. 14th International Conference on Computer Aided Verification*, volume 2404 of *LNCS*, 223–235. Springer.

Younes, H. L. S., and Simmons, R. G. 2003. VHPOP: Versatile heuristic partial order planner. *Journal of Artificial Intelligence Research* 20:405–430.

Younes, H. L. S., and Simmons, R. G. 2004a. Policy generation for continuous-time stochastic domains with concurrency. In *Proc. Fourteenth International Conference on Automated Planning and Scheduling*. AAAI Press.

Younes, H. L. S., and Simmons, R. G. 2004b. Solving generalized semi-markov decision processes using continuous phase-type distributions. In *Proc. Nineteenth National Conference on Artificial Intelligence*. AAAI Press.

Younes, H. L. S.; Kwiatkowska, M.; Norman, G.; and Parker, D. 2004. Numerical vs. statistical probabilistic model checking: An empirical study. In *Proc. 10th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, volume 2988 of *LNCS*, 46–60. Springer.

Younes, H. L. S.; Musliner, D. J.; and Simmons, R. G. 2003. A framework for planning in continuous-time stochastic domains. In *Proc. Thirteenth International Conference on Automated Planning and Scheduling*, 195–204. AAAI Press.

Younes, H. L. S. 2003. Extending PDDL to model stochastic decision processes. In *Proc. ICAPS-03 Workshop on PDDL*, 95–103.