Analysis & Tools: Stochastic Analysis

UT

May 31, 2005

AMETIST DELIVERABLE 2.4.b

Project acronym: AMETIST
Project full title: Advanced Methods for Timed Systems
Project no.: IST-2001-35304

1 Introduction & Summary

This is the Ametist third year's deliverable of Task 2.4.b on Stochastic Analysis. The main aim of this task is to achieve an integration of techniques for the analysis of timed systems with those for stochastic systems. To obtain such integration, the task has focused on extending existing, and highly successful, automated verification techniques to models that contain random aspects.

In the first two years of the Ametist project this has led to efficient algorithms for so-called fully probabilistic systems, i.e., transition systems for which the entire behaviour is probabilistic. These models, such as discrete-time and continuous-time Markov chains (DTMCs and CTMCs, respectively) are widely used in the field of performance and dependability analysis. The resulting techniques can be used to determine performance and dependability *guarantees* of Markov chains. Guarantees are described using logics. The use of logics yields an expressive framework that allows to express well-known measures, but also (new) intricate and complex performance guarantees. The power of this technique is that no matter how complex the logical guarantee, it is *automatically* checked which states in the Markov chain satisfy it. This aspect—common to verification—is new for the field of performance and dependabililty analysis where typically techniques are developed that are tailored to specific measurs-of-interest. Neither manual manipulations of Markov chains (or their high-level descriptions) are needed, nor the knowledge of any numerical technique to analyze them efficiently. This applies to any (time-homogeneous) Markov chain of any structure specified in any high-level formalism (see [5]).

In the final year of Ametist, these techniques have been refined, extended to models that exhibit randomness and non-determinism, and, moreover, more general probabilistic models have been considered for which simulation techniques are used (rather than numerical algorithms as for DTMCs and CTMCs). Emphasis has been put on the following issues:

- **Model-checking algorithms.** Algorithms have been developed for the verification of time-bounded reachability probabilities in continuous-time Markov decision processes [4], time- and cost-bounded reachability for CTMCs with both state- and impulse rewards [9], and for expectations on cost extensions of DTMCs and MDPs using discounting [11].

- **Abstraction techniques.** These techniques are focused on diminishing the state-space explosion problem as encountered in, e.g., model checking of probabilistic systems. Various strong and weak branching-time equivalences and pre-orders have been investigated and compared for discrete- and continuous-time Markov chains [6]. This is complemented by an extension of partial-order techniques for MDPs for both the linear- and branching-time perspective [2, 10] as well as a thorough investigation of branching bisimulation for MDPs [1].

- **Tools.** The ETMCC tool has been extended with features to reason and verify properties of cost extensions of DTMCs and CTMCs. This includes the algorithms that have been developed in the earlier phases of the Ametist project as well as the recent algorithms in [9]. A prototypical implementation of a tool supporting the stochastic process formalism MoDeST (see the recent overview in [16]) has been completed; this tool contains a discrete-event simulator enabling the stochastic analysis of MoDeST specifications.

- **Case studies.** In the final year of AMETIST, several case studies have been conducted and finalized: the Axxom case study on scheduling of lacquer production [8], a lightweight protocol to detect node absence in dynamic networks [7] and a reliability analysis of train radio communications [14, 12].

A more detailed summary of these contributions is given in the remainder of this report. Part of the reported work has been funded from other sources than the AMETIST project, but all of the reported work is strongly related to AMETIST in terms of relevance of results, as well as the involved researchers.

The results in the field of stochastic analysis are of very high quality and leading-edge internationally. This applies to theory developments as well as to tool developments and the practical experiments.

2 Model-checking algorithms

2.1 Time-bounded reachability in CTMDPs

A continuous-time Markov decision process (CTMDP) is a generalization of a CTMC in which both probabilistic and nondeterministic choices co-exist. [4] presents an efficient algorithm to compute the maximum (or minimum) probability to reach a set of goal states within a given time bound in a uniform CTMDP, i.e., a CTMDP in which the delay time distribution per state visit is the same for all states. The worst-case time complexity of this algorithm equals that for the fully probabilistic case except for a linear dependency on the degree of nondeterminism in thre CTMDP. It is proven that these probabilities coincide for (time-abstract) history-dependent and Markovian schedulers that resolve nondeterminism either deterministically or in a randomized way. As opposed to the discrete-time case, simple (i.e., Markovian) schedulers are shown to be sub-optimal, whereas timed schedulers are shown to outperform all other policies. ([4] is an extended version of [3].)

2.2 Discounted temporal properties

Temporal logic is two-valued: a property is either true or false. When applied to the analysis of stochastic systems, or systems with imprecise formal models, temporal logic is therefore fragile: even small changes in the model can lead to opposite truth values for a specification. [11] presents *discounted CTL (DCTL)*, a generalization of the branching-time logic CTL. DCTL achieves robustness with respect to model perturbations by giving a quantitative interpretation to predicates and logical operators. Rather than a boolean, the interpretation of a DCTL formula is a real number in the interval [0,1], where 1 corresponds to truth and 0 to falsehood. Moreover, the path operators are discounted by a discount factor in [0,1], so that weight can be given to states that are closer to the beginning of the path. [11] interprets DCTL over transition systems, DTMCs, and MDPs and presents algorithms for model checking the logic over these structures.

2.3 Time- and cost-bounded reachability in CTMCs with rewards

Markov reward models (MRMs) have been applied to the simultaneous analysis of performance and dependability of computer systems. The combination of performance and dependability is sometimes referred to as *performability*. MRMs are CTMCs with rewards assigned to their states and transitions (impulse rewards). Model checking for such MRMs has been introduced earlier along with a logic, Continuous Stochastic Reward Logic (CSRL), for the expression of relevant properties of MRMs. Existing model-checking methods have been, however, applicable to MRMs with state reward structures only. The paper [9] extends MRMs, CSRL with impulse rewards and presents new model-checking algorithms for these models. Numerical methods which are applicable for the computation of the measures defined are presented. The applicability of these algorithms is demonstrated in the context of dynamic power management.

2.4 Mapping interactive Markov chains onto CTMDPs

CTMDPs can intuitively be viewed as a common semantic model for confused generalized stochastic Petri nets (GSPNs), stochastic activity networks which are not well-defined, and for various stochastic process algebras including interactive Markov chains (IMCs). Such performance and dependability models have so far resisted any kind of analysis, as a result of the presence of nondeterminism. [13] makes the intuitive connection to CTMDPs more precise and shows how IMCs can be mapped onto CTMDPs by a sequence of sound transformations. Variants of bisimulation and trace equivalence are used to formally establish soundness of the transformations.

3 Abstraction techniques

3.1 Comparison of branching-time abstraction techniques

In the literature, a plethora of branching-time relations have been defined for fully probabilistic systems. [6] presents various semantics in the branching-time spectrum of DTMCs and CTMCs and compares them. Strong and weak bisimulation equivalence and simulation pre-orders are covered and are logically characterised in terms of the temporal logics PCTL and CSL. Apart from presenting various existing branching-time relations in a uniform manner, the contributions of this work are: (i) weak simulation for DTMCs is defined, (ii) weak bisimulation equivalence is shown to coincide with weak simulation equivalence, (iii) logical characterisation of weak (bi)simulations are provided, and(iv) a classification of branching-time relations is presented, elucidating the semantics of DTMCs, CTMCs and their interrelation.

3.2 Branching bisimulation

Instead of considering equivalences on fully probabilistic systems (like in [6]), it is interesting to consider the interplay between probabilities and non-determinism. This is considered in [1]. Branching bisimulation– a well-known notion for labeled transition systems—is considered for Hansson's alternating model as well as an equivalence notion that is based on schedulers that induce maximal probabilities. Both equivalences are shown to coincide, which allows to decide branching bisimulation by only considering a subset of schedulers. Finally, it is shown (similar to a result by van Glabbeek and Weijland in the non-probabilistic setting), that branching bisimulation coincides with coloured trace equivalence where the colours are used to code the branching potential of a system in a node.

3.3 Partial-order reduction

Partial order reduction (POR) has been used to alleviate the state explosion problem in model checkers for nondeterministic systems. The method relies on exploring only a fragment of the full state space of a program that is enough to assess the validity of a property. For both linear- and branching-time specifications, methods have been developed to apply POR to model checking. Results in [10] give criteria on applying POR to verify quantitative LTL properties for MDPs. [2] presents POR criteria branching-time properties of MDPs, such as PCTL and compares these criteria to those for linear-time properties.

3.4 Divergence

When a process is capable of executing an unbounded number of non-observable actions it is said to be *divergent*. Different capabilities of an observer to identify this phenomen along the execution leads to different divergent sensitive semantics. [] develops sound and complete axiomatisations for the divergence sensitive spectrum of weak bisimulation equivalence. The axiomatisations separates the axioms concerning recursion and those that capture the essence of diverging behaviour.

(Although this paper is based on transition systems, its motivation comes from stochastic models where instantaneous and stochastically delayed actions co-exist.)

4 Tools

In order to facilitate the analysis of Modest models, we have developed the prototype tool Motor [16]. Modest is a very expressive language, covering a wide range of timed, probabilistic, nondeterministic, and stochastic models. The spectrum of covered models includes ordinary labeled transition systems, discrete and continuous time Markov chains and decision processes, generalized semi-Markov processes, and timed and probabilistic timed automata. These submodels play a crucial role in the context of Motor. The enormous expressiveness of Modest implies that no generic analysis algorithm is at hand. Instead, Motor aims at supporting a variety of analysis algorithms tailored to the variety of analyzable submodels. The philosophy behind Motor is to connect Modest to existing tools, rather than re-implementing existing analysis algorithms anew. In joint efforts with the Moebius developers, MoToR is linked to the powerful solution techniques of Möbius for quantitative assessment. The main objective was to simulate Modest models by means of the Möbius distributed discrete-event simulator, because a stochastic simulator can cope with one of the largest class of models expressible in Modest.

The current version of MoToR is available (open source) from the web.

5 Case studies

The case studies that have been conducted in the 3rd year of Ametist have focused on applying the techniques and tools that have been developed during the first two years, focusing on stochastic aspects. The variety of the case studies (scheduling, load balancing and reliability analysis) indicates that the developed algorithms and tools are not tailored to a specific application area, but are rather widely applicable. All reported case studies use the modeling formalism MoDeST and its accompanying tool MoToR.

5.1 Stochastic assessment of schedules for lacquer production

The effect of faulty behavior on a hard real-time scheduling problem from the domain of lacquer production is investigated in [8]. The scheduling problem is first solved using the timed model-checker UPPAAL. The resulting schedules are then embedded in a MoDeST failure model of the lacquer production line, and analyzed with the discrete event simulator of Mobius. This approach allows one to assess the quality of the schedules with respect to timeliness, utilization of resources, and sensitivity to different assumptions about the reliability of the production line.

5.2 Monitoring node absence in self-configuring networks

[7] is concerned with the design of a distributed algorithm to monitor the availability of nodes in self-configuring networks. It is intended as extension to service discovery protocols such as SSDP, SLP, Rendezvous and Jini that allow for fast node detection. The simple scheme to regularly probe a node—“are you still there?”—may easily lead to over- or underloading. The essence of our algorithm is therefore to automatically adapt the probing frequency. It is shown that a self-adaptive scheme to control the probe load, originally proposed as an extension to the UPnP (Universal Plug and Play) standard, leads to an unfair treatment of nodes: some nodes probe fast while others almost starve. An alternative distributed algorithm is proposed that overcomes this problem and that tolerates highly dynamic network topology changes. The algorithm is very simple and can be implemented on large networks of small computing devices such as mobile phones, PDAs, and so on. The distributed algorithms are modeled using MoDeST and are analyzed by means of the discrete-event simulator of the MoToR tool.

5.3 Reliability analysis of train radio communications

In this case study (see [14]) StoCharts, a stohastic extension of UML statecharts is used to model a part of the European Train Control System (ECTS) specification, focusing on the risks of wireless communication failures in future high-speed cross-European trains. Stochastic model checking with the model checker Prover enables us to derive constraints under which the central quality requirements are satisfied by the model. This case study illustrates the flexibility and maturity of StoCharts to model real problems in safety critical system design. The cornerstones of a mechanizable translation of StoCharts to MoDeST are detailed out in [12]. An informal introduction to StoCharts is provided in [15].

References

[1] S. Andova and T. Willemse, *Equivalences for silent transitions in probabilistic systems (extended abstract*, Electronic Notes in Theoretical Computer Science **2** (2005), no. 128, 53–66.

[2] C. Baier, P.R. D'Argenio, and M. Größer, *Partial order reducction for probabilistic branching time*, 3rd Workshop of Quantitative Aspects of Programming Languages, *QAPL'05*, 2005.

[3] C. Baier, B. Haverkort, H. Hermanns, and J.-P. Katoen, *Efficient computation of time-bounded reachability probabilities in uniform continuous-time Markov decision processes*, Tools and Algorithms for the Construction and Analysis of Systems (TACAS) (Barcelona, Spain), Lecture Notes in Computer Science, Springer-Verlag, 2004.

[4] ———, *Efficient computation of time-bounded reachability probabilities in uniform continuous-time Markov decision processes*, Theoretical Computer Science (2005), Accecpted for publication.

[5] ———, *Model checking meets performance evaluation*, ACM SIGMETRICS Performance Evaluation Review **32** (2005), 10–15.

[6] C. Baier, J.-P. Katoen, H. Hermanns, and Verena Wolf, *Comparative branching-time semantics for Markov chains*, Information and Computation (2005), Accepted, publication date unknown.

[7] H. Bohnenkamp, J. Gorter, J. Guidi, and J.-P. Katoen, *Are you still there? a lightweight algorithm to monitor node absence in self-configuring networks*, Dependable Systems and Networks (DSN) (2005), To appear, 6 pages.

[8] H.C. Bohnenkamp, H. Hermanns, R. Klaren, A. Mader, and Y.S. Usenko, *Synthesis and stochastic assessment of schedules for lacquer production*, Proc. 1st Int. Conf. on Quantitative Evaluation of Systems (QEST), IEEE CS Press, sep 2004, pp. 28–37.

[9] L. Cloth, J.-P. Katoen, M. Khattri, and R. Pulungan, *Model checking Markov reward models with impulse rewards*, Dependable Systems and Networks (DSN) (2005), To appear, 10 pages.

[10] P.R. D'Argenio and P. Niebert, *Partial order reduction for concurrent probabilistic programs*, Proc. of the 1st International Conference on Quantitative Evaluation of Systems *2004* (B.R. Haverkort et al., ed.), IEEE Computer Society Press, 2004, pp. 240–249.

[11] L. de Alfaro, M. Faella, T. A. Henzinger, R. Majumdar, and M.I.A. Stoelinga, *Model checking discounted temporal properties*, Theoretical Computer Science (2005), Accecpted for publication.

[12] H. Hermanns, D.N. Jansen, and Y.S. Usenko, *From stocharts to modest: a comparative reliability analysis of train radio communications*, Proc. 5th International Workshop on Software and Performance (WOSP'05) (Palma de Mallorca, Spain), July 2005, To appear. Also appeared as a SFB/TR 14 AVACS Technical Report ATR-002.

[13] H. Hermanns and S. Johr, *Towards analysing nondeterministic specifications of stochastic systems*, submitted to FORMATS 2005.

[14] D.N. Jansen and H. Hermanns, *Dependability checking with stocharts: Is train radio reliable enough for trains?*, In 1st International Conference on Quantitative Evaluation of Systems (QEST), IEEE CS Press, 2004, pp. 250–259.

[15] ———, *QoS modelling and analysis with UML-statecharts: the StoCharts approach*, ACM SIGMETRICS Performance Evaluation Review **32** (2005), 28–33.

[16] J.-P. Katoen, Henrik Bohnenkamp, H. Hermanns, and J. Klaren, *Embedded software analysis with MOTOR*, LNCS, pp. 268–294, Springer, Bertinoro, Italy, 2004.