Symbolic Model Checking of Concurrent Probabilistic Systems Using MTBDDs and the Simplex Algorithm

Marta Kwiatkowska[1⋆], Gethin Norman[1⋆], David Parker[1], and Roberto Segala[2⋆⋆]

[1] University of Birmingham, Birmingham B15 2TT, United Kingdom
{M.Z.Kwiatkowska,G.Norman,D.A.Parker}@cs.bham.ac.uk
[2] Dipartimento di Scienze dell'Informazione, Università di Bologna,
Mura Anteo Zamboni 7, 40127 Bologna, Italy
segala@cs.unibo.it

Corresponding author: Marta Kwiatkowska
Tel +44 121 414 7264, FAX +44 121 4281

Abstract. Symbolic model checking for purely probabilistic processes using MTBDDs [12] was introduced in [4] and further developed in [20, 3]. In this paper we consider models for *concurrent* probabilistic systems similar to those of [28, 7, 5] and the concurrent Markov chains of [35, 13], which extend the purely probabilistic processes through the addition of nondeterministic choice. As a specification formalism we use the probabilistic branching-time temporal logic PBTL of [5, 7], which allows us to express properties such as "under any scheduling of nondeterministic choices, the probability of ϕ holding until ψ is true is at least 0.78". In [5, 7] it is shown that the verification of "until" properties can be reduced to a linear programming problem and solved with the help of e.g. the simplex algorithm, but no symbolic model checking is considered. Based on the algorithms of [5, 7], we derive *symbolic* model checking procedure for PBTL over concurrent probabilistic systems using MTBDDs, and extend it with fairness constraints. We furthermore implement an experimental model checker using the Colorado University Decision Diagram (CUDD) package [32]. Our key contribution is an implementation of the *simplex algorithm* in terms of MTBDDs.

Keywords: temporal logic, symbolic model checking, probabilistic verification, concurrent Markov chains.

1 Introduction

There have been many advances in the BDD technology since BDDs were first introduced [9] and applied to symbolic model checking [10, 24]. There are now several free and commercial BDD packages in existence, as well as a range of alternative techniques for improving efficiency. Model checking tools (to mention `smv`, `SPIN`, `fdr2`) are extensively used by industrial companies in the process of developing new designs for e.g. hardware circuits, network protocols, etc. There have also been encouraging developments in model checking of real-time and hybrid systems [6, 8].

One area that is lagging behind as far as experimental work is concerned, despite the fact that the fundamental verification algorithms have been known for over a decade [35, 19, 25], is model checking of *probabilistic* systems. This is particularly unsatisfactory since many systems currently being designed would benefit from probabilistic analysis performed *in addition to* the conventional, qualitative, checks involving temporal logic formulas or reachability analysis available in established model checking tools. This includes not only quality of service properties such as "with probability 0.9 or greater, the system will respond to the request within time t", but also performance analysis, currently separate from model

⋆ supported in part by EPSRC grant GR/M04617
⋆⋆ supported in part by EPSRC grant GR/M13046

checking, which allows the computation of characteristics such as long-run average, resource utilization, etc.

In order to support the verification of probabilistic systems, BDD-based packages must allow a compact representation for sparse probability matrices. Such a representation, Multi-Terminal Binary Decision Diagrams (MTBDDs), was proposed in [12] along with matrix algorithms. Based on [18], an MTBDD-based *symbolic* model checking procedure for purely probabilistic processes (state-labelled discrete Markov chains) for the logic PCTL of [18] (a probabilistic variant of CTL which includes formulas such as $[\phi\ \mathcal{U}\ \psi]_{\geq\lambda}$ meaning the probability of satisfying $\phi\ \mathcal{U}\ \psi$ is at least λ), was first presented in [4]. The algorithm for checking $[\phi\ \mathcal{U}\ \psi]_{\geq\lambda}$ involves computing the probability measure (assuming the usual measure on the σ-algebra of sets of infinite paths induced from the cones of finite paths) of the set of all paths satisfying $\phi\ \mathcal{U}\ \psi$, and reduces to solving a system of *linear equations*. The bounded "until" variant of this algorithm was implemented within Probabilistic Verus [20, 3], a model checker for synchronous parallel randomized programs. The research concerning Probabilistic Verus focuses on semantics and case studies, featuring timing and performance analysis, but does not include support for nondeterminism.

In this paper we consider models for *concurrent probabilistic systems* similar to those of [28, 7, 5] and slightly more general than the concurrent Markov chains of [35, 13]. These admit *nondeterministic choice* between discrete probability distributions on the successor states, and are particularly appropriate for the representation of randomized distributed algorithms, fault-tolerant and self-stabilising systems. The presence of nondeterminism means that no single probability space can be associated with the computation paths of our models: the probability space is determined only for a given *adversary* (or *scheduler*) which selects one of possibly several probability distributions available in each state. Consequently, the logic we define for concurrent probabilistic systems combines features of the branching-time logic CTL with probabilistic reasoning as follows: two variants of "until" are introduced, $\exists\mathcal{U}$ and $\forall\,\mathcal{U}$, with quantification defined over all adversaries, respectively all *fair* adversaries, and furthermore the probability measure is calculated similarly to the case for PCTL above once an adversary has been specified. Fairness is needed for the verification of liveness properties. The logic we use is Probabilistic Branching-Time Temporal Logic (PBTL) of [5] (it is essentially equivalent to pCTL in [7]). The model checking procedure, already proposed in [5, 7][1] albeit not in conjunction with a symbolic method, reduces to the computation of the *minimum/maximum* probability over the set of simple adversaries, as opposed to exact probability calculation possible in the case of PCTL. Alternatively, we can derive a set of *linear inequalities*, and maximize/minimize the sum of the components of the solution vector subject to the constraints given by the inequalities. As suggested in [7], this linear programming problem can be solved with the help of e.g. the simplex algorithm, see e.g. [27].

The main contribution of this paper is a *symbolic model checking* method for PBTL over concurrent probabilistic systems using MTBDDs, which incorporates optional *fairness constraints*, together with an implementation of an MTBDD-based experimental tool. Though several BDD-based packages are in existence, few support MTBDDs as well as fully fledged matrix manipulation algorithms. The Colorado University Decision Diagram (CUDD) package of Fabio Somenzi [32] is one example of the latter type, which has already been used for probability calculations [1, 16]; we therefore adopt it for our experiments. Our tool inputs a sparse matrix description of the system, a PBTL formula and the type of satisfaction requested (with or without fairness), and outputs the vector of states for which this formula holds with respect to the chosen satisfaction relation. The model checking algorithm is driven by the syntax tree of the formula, and proceeds in a bottom-up fashion making use of both the BDD operations (in the case of e.g. Boolean connectives), as well MTBDD-based

[1] The model checking algorithms of [5, 7] coincide for the case of verification without fairness constraints. Fairness is not considered in [7], but see the more recent probabilistic fairness of [14].

probability calculations. We defer the design of a system description language till later, concentrating instead on the model checking aspects. The minimum/maximum probability calculation for "until" formulas is implemented in two ways. The first is via the iterative refinement (known to be better for numerical stability, see e.g. [22], but potentially leading to slower running times if the probability function converges slowly), and the other using *simplex* adapted to MTBDDs.

The use of simplex was already suggested in [7]. Since no symbolic model checking was considered there, this implied the need for explicit state space construction or conversion to the internal representation of some linear algebra package. Coding simplex in terms of MTBDDs avoids unnecessary conversion, and thus affords a single, compact representation of the system. Since the model checking algorithm is a combination of BDD and MTBDD operations, this allows us to use the BDD operations at maximum efficiency, and yields the additional advantage that conventional, qualitative, BDD-based symbolic checks can be easily combined together with probabilistic analysis within the same tool. The reason we use simplex in preference to e.g. the exterior point ellipsoid method is that it offers the best performance in practical cases [23, 27]; an alternative algorithm that might also be worth investigating is the interior point algorithm due to Karmarkar [23] (also known as the interior point ellipsoid method). Though exponential in the worst case, the examples that exhibit exponential time are considered artificial and are rarely found in practical applications [27]. It has been shown that, for a certain natural probabilistic space, simplex is time polynomial on average [27]. Combined with BDDs, the simplex algorithm does not directly benefit from their recursive structure as it relies on accessing individual rows and columns, but has the potential to benefit from the compactness of the representation, though its performance compared to sparse matrices is difficult to predict. We have compared the MTBDD implementation with sparse matrices using MATLAB, indeed with encouraging results. To the authors' knowledge, this is the first MTBDD-based implementation of a linear programming problem; we intend to learn from this experience and, if proved efficient in practice, this opens up the possibility of coding other linear algebra operations in terms of MTBDDs, thus making it possible for BDD-based model checkers to incorporate performance analysis.

Related work: Much has been published concerning probabilistic logics and verification, see e.g. [35, 13, 25, 26, 18, 17, 7, 5], but only a few papers deal with BDD-based approaches to probability calculation. Probabilistic Verus [20, 3] offers an extensive C-like system description language, but does not permit nondeterminism. We know of three further approaches based on BDD technology: Algebraic Decision Diagrams (ADDs) [1, 16], Factored Edge-Valued BDDs (FEVBDDs) [34], and Decision Node BDDs (DNBDDs) [30, 31]. Of these, only the latter admits action choice in addition to probabilistic choice; all three significantly differ from the discussion here in that they focus on the calculation of stationary-state probabilities instead of model checking of probabilistic temporal logic formulas. The ADDs of [1, 16] coincide as data structures with MTBDDs (henceforth we shall not distinguish between the two). In [1] three matrix multiplication algorithms for MTBDDs are described which differ in terms of performance depending on the structure of the matrices. In [34] FEVBDDs are used to derive an alternative representation of matrices, together with matrix operations, and are applied to Integer Linear Programming, resulting in space efficient but slow algorithms based on infinite precision arithmetic. It is not clear how the FEVBDD representation can be made to work with floating point numbers, used as standard in numerical probabilistic analysis. DNBDDs [30, 31] offer a data structure for representing stochastic transition diagrams efficiently, but we are not aware of experiments with probability calculations for DNBDDs.

2 Concurrent Probabilistic Systems

In this section, we briefly summarise our underlying model for concurrent probabilistic systems; the reader is referred to [5] for more details. Our model is based on "Markov decision

processes"; it slightly generalises "Concurrent Markov Chains" of [35,13] and essentially coincides with "simple deterministic automata" of [28]. Some familiarity with Markov chains and probability theory is assumed, see e.g. [33].

We begin by recalling the standard definition of (discrete) Markov chains and the associated probability measure. A *Markov chain* is a pair $M = (S, \mathbf{P})$ where S is a finite set of states and $\mathbf{P} : S \times S \to [0,1]$ a *transition matrix*, i.e. $\sum_{t \in S} \mathbf{P}(s,t) = 1$ for all $s \in S$. A *path* in M is a nonempty (finite or infinite) sequence $\pi = s_0 s_1 s_2, \ldots$ where s_i are states and $\mathbf{P}(s_{i-1}, s_i) > 0$, $i \in \mathbb{N}$. The length of a path, denoted $|\pi|$, is defined in the usual way. A path π is called a *full path* iff it is infinite. We model terminating behaviour by repeating the final state infinitely often. The first state of π is denoted by $\mathit{first}(\pi)$, and the last by $\mathit{last}(\pi)$ (if it exists).

Path_{ful} denotes the set of full paths, whereas $\mathit{Path}_{ful}(s)$ is the set of full paths π with $\mathit{first}(\pi) = s$. For $s \in S$, let $\mathcal{F}(s)$ be the smallest σ-algebra on $\mathit{Path}_{ful}(s)$ which contains the basic cylinders $\{\pi \in \mathit{Path}_{ful} : \rho \text{ is a prefix of } \pi\}$ where ρ ranges over all finite paths starting in s. The probability measure Prob on $\mathcal{F}(s)$ is the unique measure with $\mathit{Prob}\{\, \pi \in \mathit{Path}_{ful}(s) : \rho \text{ is a prefix of } \pi \,\} = \mathbf{P}(\rho)$ where $\mathbf{P}(s_0 s_1 \ldots s_k) = \mathbf{P}(s_0, s_1) \cdot \mathbf{P}(s_1, s_2) \cdot \ldots \cdot \mathbf{P}(s_{k-1}, s_k)$.

Concurrent probabilistic systems generalise ordinary Markov chains in that they allow possibly several probability distributions to be enabled in a given state. The selection of a probability distribution is made by an adversary (also known as a scheduler), and then the process makes a transition to a state determined by the selected distribution. We denote the set of discrete probability distributions over a set S by $\mu(S)$. Therefore, each $p \in \mu(S)$ is a function $p : S \to [0,1]$ such that $\sum_{s \in S} p(s) = 1$.

Definition 1 (Concurrent Probabilistic System). *A* concurrent probabilistic system *is a pair $\mathcal{S} = (S, \mathit{Steps})$ where S is a finite set of states and Steps a function which assigns to each state $s \in S$ a finite, non-empty set $\mathit{Steps}(s)$ of distributions on S. Elements of $\mathit{Steps}(s)$ are called* transitions.

Paths in a concurrent probabilistic system arise by resolving both the nondeterministic and probabilistic choices. A *path* of the system $\mathcal{S} = (S, \mathit{Steps})$ is a non-empty finite or infinite sequence $\pi = s_0 \xrightarrow{p_0} s_1 \xrightarrow{p_1} s_2 \xrightarrow{p_2} \cdots$ where $s_i \in S$, $p_i \in \mathit{Steps}(s_i)$ with $p_i(s_{i+1}) > 0$ for all $0 \le i \le |\pi|$.

We now introduce *adversaries* of concurrent probabilistic systems.

Definition 2 (Adversary of a Concurrent Probabilistic System). *An* adversary *(or scheduler) of a concurrent probabilistic system $\mathcal{S} = (S, \mathit{Steps})$ is a function A mapping every finite path ρ of $\mathcal{S}$ to a distribution $A(\rho)$ on S such that $A(\rho) \in \mathit{Steps}(\mathit{last}(\rho))$ is a transition in $\mathcal{S}$. We use $\mathbf{A}(\mathcal{S})$ to denote the set of all adversaries of $\mathcal{S}$.*

Note that an adversary can choose a different distribution every time the system returns to the same state, which means that there are infinitely many adversaries even if the system has finitely many states. An adversary A of $\mathcal{S}$ is called *simple* iff for every state $s \in S$ there exists a transition $\mu_s \in \mathit{Steps}(s)$ with $A(\rho) = \mu_{\mathit{last}(\rho)}$ where ρ ranges over all finite paths. It turns out that for model checking it suffices to consider only the simple adversaries[2] which resolve the nondeterminism by selecting the same distribution every time a given state is reached – independent of the past history; for more detail see [5].

For an adversary A of a concurrent probabilistic system $\mathcal{S} = (S, \mathit{Steps})$ we define Path^A_{fin} to be the set of finite paths such that $\mathit{step}(\rho, i) = A(\rho^{(i)})$ for all $1 \le i \le |\rho|$, and Path^A_{ful} to be the set of paths in Path_{ful} such that $\mathit{step}(\rho, i) = A(\rho^{(i)})$ for all $i \in \mathbb{N}$, where $\rho(i)$ denotes the i-th state of ρ, $\mathit{step}(\rho, i)$ is the distribution selected by the i-th step, and $\rho^{(i)}$ is the i-th prefix of ρ.

[2] These correspond to "stationary policies" in Markov decision processes.

With each adversary we associate a sequential Markov chain, together with the induced probability measure $Prob$, which can be viewed as a set of paths in $\mathcal{S}$. Formally, if A is an adversary of the concurrent probabilistic system $\mathcal{S}$, then $MC^A = (Path^A_{fin}, \mathbf{P}^A)$ is a Markov chain where:

$$\mathbf{P}^A(\rho, \rho') = \begin{cases} p(s) \text{ if } A(\rho) = p \text{ and } \rho' = \rho \xrightarrow{p} s \\ 0 \quad \text{otherwise} \end{cases}$$

Since we allow nondeterministic choice between probability distributions, we may have to impose *fairness constraints* in order to ensure that certain liveness properties can be verified. In a distributed environment fairness corresponds to a requirement for each each concurrent component to progress whenever possible. Without fairness, certain liveness properties may trivially fail to hold in the presence of simultaneously enabled transitions of a concurrent component.

We define a fulpath π of a concurrent probabilistic system to be *fair* iff, whenever a state s is visited infinitely often in π, each nondeterministic alternative which is enabled in s is taken infinitely often in π. More formally, we have the definition below.

Definition 3 (Fairness of Fulpaths). *Let $\mathcal{S} = (S, Steps)$ be a concurrent probabilistic system and π a fulpath in $\mathcal{S}$. π is called* fair *iff, for each $s \in inf(\pi)$ and each $\mu \in Steps(s)$, there are infinitely many indices i with $\pi(i) = s$ and $step(\pi, i) = \mu$. $Fair(\mathcal{S})$, abbreviated by $Fair$, denotes the set of fair fulpaths in $\mathcal{S}$.*

Now an adversary is *fair* if any choice of transitions that becomes enabled infinitely often along a computation path is taken infinitely often.

Definition 4 (Fair Adversaries). *Let $\mathcal{S} = (S, Steps)$ be a concurrent probabilistic system and F an adversary for $\mathcal{S}$. F is called* fair *iff $Prob(Fair^F(s)) = 1$ for all $s \in S$. We let $\mathbf{A}_{fair}(\mathcal{S})$ denote the set of fair adversaries.*

The interested reader is referred to [14, 15] where the above fairness notion is termed 'path fairness', and an alternative, 'probabilistic fairness', which coincides with our notion for finite state systems and a large class of properties, is introduced.

Strict fairness of [5] can also be considered. We note that differences between strict fairness and the above fairness are subtle, and the model checking algorithms are more involved in the case of the former notion.

3 The Logic PBTL and Its Model Checking Procedure

In this section, based on [5, 7], we recall the syntax and semantics of the probabilistic branching-time temporal logic PBTL and outline the corresponding model checking procedure. PBTL combines features of CTL [11] and PCTL [18]. In particular, the temporal operator $\mathcal{U}$ ("until") and the path quantifiers $\forall$ and $\exists$ are taken from CTL, and the probabilistic operator $[\cdots]_{\geq\lambda}$ from PCTL.

Let AP denote a finite set of atomic propositions. A *PBTL structure* is a tuple $(\mathcal{S}, \mathrm{AP}, L)$ where $\mathcal{S} = (S, Steps)$ is a concurrent probabilistic system and $L : S \to 2^{\mathrm{AP}}$ is a labelling function which assigns to each state $s \in S$ a set of atomic propositions. By abuse of notation we shall refer to the PBTL structure $(\mathcal{S}, \mathrm{A}P, L)$ over the concurrent probabilistic system $\mathcal{S} = (S, Steps)$ as $\mathcal{S}$.

The syntax of PBTL is given below. To simplify this presentation we omit the "bounded until" and "next state" operators which can be easily added.

Definition 5 (Syntax of PBTL**).** *The syntax of* PBTL *formulas is defined as follows:*

$$\phi ::= \texttt{true} \mid a \mid \phi_1 \wedge \phi_2 \mid \neg\phi \mid [\, \phi_1 \; \exists\mathcal{U} \; \phi_2 \,]_{\sqsupseteq\lambda} \mid [\, \phi_1 \; \forall\mathcal{U} \; \phi_2 \,]_{\sqsupseteq\lambda}$$

where a is an atomic proposition, $\lambda \in [0, 1]$, and $\sqsupseteq$ is either $\geq$ or $>$. Any formula of the form $\phi_1 \mathcal{U} \phi_2$, where ϕ_1, ϕ_2 are PBTL formulas, is defined to be a path formula.

PBTL formulas are interpreted over states of concurrent probabilistic systems, whereas path formulas over paths. The branching time quantifiers $\exists$ and $\forall$ involve quantification over adversaries: $\exists$ means "there exists an adversary" of a given type, and $\forall$ "for all adversaries" of a given type.

Definition 6 (Satisfaction Relation for PBTL**).** *Given a PBTL-structure* $(\mathcal{S}, \mathrm{AP}, L)$*, a set Adv of adversaries of* $\mathcal{S}$*, and* PBTL *formula* ϕ*, we define the satisfaction relation* $s \models_{Adv} \phi$ *inductively as follows:*

$$
\begin{array}{ll}
s \models_{Adv} \mathtt{true} & \text{for all } s \in S \\
s \models_{Adv} a & \Leftrightarrow a \in L(s) \\
s \models_{Adv} \phi_1 \wedge \phi_2 & \Leftrightarrow s \models_{Adv} \phi_1 \text{ and } s \models_{Adv} \phi_2 \\
s \models_{Adv} \neg\phi & \Leftrightarrow s \not\models_{Adv} \phi \\
s \models_{Adv} [\phi_1 \; \exists\mathcal{U} \; \phi_2]_{\sqsupseteq\lambda} & \Leftrightarrow \mathit{Prob}(\{\pi \mid \pi \in \mathit{Path}^A_{ful}(s) \;\&\; \pi \models_{Adv} \phi_1 \,\mathcal{U}\, \phi_2\}) \sqsupseteq \lambda \\
 & \quad \text{for some adversary } A \in \mathit{Adv} \\
s \models_{Adv} [\phi_1 \; \forall\mathcal{U} \; \phi_2]_{\sqsupseteq\lambda} & \Leftrightarrow \mathit{Prob}(\{\pi \mid \pi \in \mathit{Path}^A_{ful}(s) \;\&\; \pi \models_{Adv} \phi_1 \,\mathcal{U}\, \phi_2\}) \sqsupseteq \lambda \\
 & \quad \text{for all adversaries } A \in \mathit{Adv} \\
 & \\
\pi \models_{Adv} \phi_1 \,\mathcal{U}\, \phi_2 & \Leftrightarrow \text{there exists } k \geq 0 \text{ such that } \pi(k) \models_{Adv} \phi_2 \\
 & \quad \text{and for all } j = 0, 1, \ldots, k-1 \; \pi(j) \models_{Adv} \phi_1
\end{array}
$$

We denote $\models_{\mathbf{A}(\mathcal{S})}$ *by* $\models$ *(satisfaction for* all *adversaries) and* $\models_{\mathbf{A}_{fair}(\mathcal{S})}$ *by* $\models_{fair}$ *(satisfaction for* all fair *adversaries).*

The above definition in fact gives rise to an indexed family of satisfaction relations $\models_{Adv}$ [5], of which we only consider two, $\models$ and $\models_{fair}$. If fairness is not considered, there are obvious similarities between the logic PBTL and the probabilistic logics in the extant literature. The difference with the logic PCTL of [29] is that we label the states with atomic propositions, whereas in [29] action-labelled concurrent probabilistic systems are considered. As already mentioned, the logic pCTL of [7] essentially agrees with PBTL except that pCTL uses a probabilistic operator $\mathbb{P}_*(\cdot)$ for formulas containing the until operator. One can translate formulas of pCTL into PBTL formulas, and vice versa. For example the PBTL formula $[\, \phi_1 \; \forall\mathcal{U} \; \phi_2 \,]_{\sqsupseteq\lambda}$ can be identified with the pCTL formula $\mathbb{P}_{\sqsupseteq\lambda}(\phi_1 \mathcal{U} \phi_2)$, while $[\, \phi_1 \; \exists\mathcal{U} \; \phi_2 \,]_{>\lambda}$ corresponds to $\neg\mathbb{P}_{\leq\lambda}(\phi_1 \mathcal{U} \phi_2)$.

3.1 Model Checking for the Satisfaction Relation $\models$

With the exception of "until" formulas and fairness, model checking for PBTL is straightforward, see [7, 5]. First consider the case of $\models$ without fairness and universal "until" formulas. To verify $[\phi \; \forall\mathcal{U} \; \psi]_{\sqsupseteq\lambda}$ we first compute the sets of states S_ϕ, S_ψ that satisfy ϕ and ψ respectively. Then we calculate the *minimum probability*:

$$p^{\min}_s(\phi \;\mathcal{U}\; \psi) = \inf\{p^A_s(\phi \;\mathcal{U}\; \psi) \mid A \in \mathbf{A}(\mathcal{S})\}$$

where $p^A_s(\phi \;\mathcal{U}\; \psi) = \mathit{Prob}(\{\pi \mid \pi \in \mathit{Path}^A_{ful}(s) \;\&\; \pi \models \phi \;\mathcal{U}\; \psi\})$ and establish whether the result is $\sqsupseteq \lambda$. This can be achieved by one of the methods below.

Let S_ϕ, S_ψ be as before, $S^{\mathrm{yes}} = S_\psi$ (the set of all states which trivially satisfy $\phi \;\mathcal{U}\; \psi$ with probability 1), and let $S^{>0}$ denote the set of states from which, for all adversaries, one can reach a state in S_ψ via a path through the states in S_ϕ with positive probability (this can be computed by standard BDD-based reachability analysis). Let $S^{\mathrm{no}} = S \setminus S^{>0}$ (states which definitely do not satisfy $\phi \;\mathcal{U}\; \psi$, and must hence be mapped to 0), and $S^{?} = S^{>0} \setminus S_\psi$ (states which satisfy $\phi \;\mathcal{U}\; \psi$ with some, yet unknown, positive probability. We can establish satisfaction as follows.

The iterative method: The minimum probability $p^{\min}_s(\phi \;\mathcal{U}\; \psi)$ can be characterized as the limit, where n tends to ∞, of the following approximations $p^{\min}_{s,n}(\phi \;\mathcal{U}\; \psi)$, and hence can be

approximated by iteration [5, 2]. Put $p^{\min}_{s,n}(\phi\,\mathcal{U}\,\psi) = 1$ if $s \in S_\psi (= S^{\text{yes}})$, and $p^{\min}_{s,n}(\phi\,\mathcal{U}\,\psi) = 0$ if $s \in S \setminus (S_\phi \cup S_\psi)$, for $n = 0, 1, 2, \cdots$. For $s \in S_\phi \setminus S_\psi$ define $p^{\min}_{s,0}(\phi\ \mathcal{U}\ \psi) = 0$ and for $n = 0, 1, 2, \cdots$ put

$$p^{\min}_{s,n+1}(\phi\ \mathcal{U}\ \psi) = \min\{\Sigma_{t\in S}\ \mu(t) \cdot p^{\min}_{t,n}(\phi\ \mathcal{U}\ \psi) \mid \mu \in \mathit{Steps}(s)\}$$

In the summation $\Sigma_{t\in S}$ above it is sufficient to sum up over the set $S^{>0}$ instead of S.

The linear optimization method: Alternatively, the values $p^{\min}_s(\phi\ \mathcal{U}\ \psi)$ can also be computed by solving linear optimization problems [7, 13, 2]. We define $x_s = 1$ if $s \in S_\psi (= S^{\text{yes}})$ and $x_s = 0$ if $s \in S^{\text{no}}$. For the remaining states $s \in S^?$, the vector $p^{\min}_s(\phi\ \mathcal{U}\ \psi)$ is the solution of the linear maximization problem $0 \le x_s \le 1$ and

$$x_s \le \Sigma_{t\in S}\ \mu(t) \cdot x_t,\ \mu \in \mathit{Steps}(s)$$

where $\Sigma_{s\in S^?} x_s$ is maximal.

For the case of *existential* "until" formulas $[\phi\ \exists\mathcal{U}\ \psi]_{\sqsupseteq\lambda}$ we calculate the maximum probability $p^{\max}_s(\phi\mathcal{U}\psi)$, which can be defined similarly [5, 2], either using an iterative method similar to the above or by reduction to the dual linear programming problem (see Subsection 5.3).

3.2 Model Checking with Fairness

In contrast to the approach of CTL, where model checking with respect to the satisfaction relation with fairness can be reduced to satisfiability checking for ordinary satisfaction, but for an implication of the form $\phi_{\text{fair}} \to \phi$ where ϕ_{fair} is the fairness constraint expressed as a CTL* formula, and ϕ is the CTL formula to model check, a similar approach does not apply for PBTL. This is because of a mismatch between quantification needed for $\models$ (over all adversaries) and that needed for $\models_{fair}$ (over *all* adversaries).

Nevertheless, model checking for $\models_{fair}$ reduces to that for $\models$ on the basis of the following statement. Part (1) of this statement is due to [5], and part (2) is an improvement of [5] shown in [2].

Theorem 1. *1.* $s \models_{fair} [\phi\ \exists\mathcal{U}\ \psi]_{\sqsupseteq\lambda}$ *iff* $p^{\max}_s(\phi\ \mathcal{U}\ \psi) \sqsupseteq \lambda$

2. $s \models_{fair} [\phi\ \forall\mathcal{U}\ \psi]_{\sqsupseteq\lambda}$ *iff* $1 - p^{\max}_s(a^?\ \mathcal{U}\ \neg a^+) \sqsupseteq \lambda$ *where* $a^+, a^?$ *are new atomic propositions, with* a^+ *denoting the set of states from which, for some adversary, one can reach a state in* S_ψ *via a path through states in* S_ϕ *with positive probability (computed by standard BDD-based reachability analysis), and* $a^? = a^+ \setminus S_\psi$.

Example: Let $\mathcal{S}$ be the PBTL-structure of Figure 1 and $\phi = [a\ \forall\mathcal{U}\ b]_{>\frac{1}{4}}$. The linear programming problem that ϕ gives rise to is also given in Figure 1. It yields for the case with no fairness $x_s = \frac{1}{4}, x_r = 0$ and $x_t = \frac{5}{8}$ as optimal in two iterations of simplex. For the case with fairness one obtains $x_s = \frac{1}{4}, x_r = \frac{5}{8}$ and $x_t = \frac{5}{8}$ in three iterations.

4 Representing Concurrent Probabilistic Systems with MTBDDs

In [12] it is shown how one can generalize BDDs to represent matrices in terms of so-called *multi-terminal* binary decision diagrams (MTBDDs). One can think of MTBDDs as function graphs over Boolean variables similar to BDDs, except that they allow finitely many possibly real values as the terminal nodes, instead of just 0 and 1. For the definition of MTBDDs see [12], and BDDs see e.g. [11].

Let D be a finite set of values, possibly reals. D–valued matrices can be represented by MTBDDs as follows. Consider a square $2^m \times 2^m$–matrix A (non-square matrices can be dealt with by padding with all-zero columns or rows). Its elements a_{ij} can be viewed as

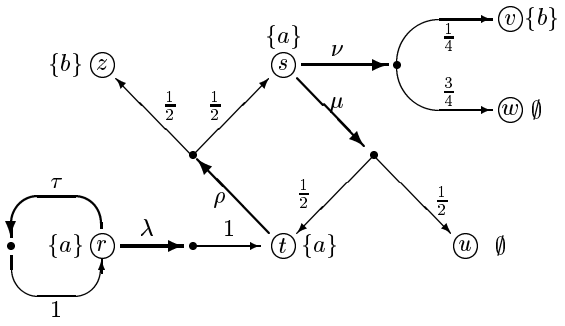

Maximize

$$x_r + x_s + x_t$$

subject to the constraints:

$$\begin{array}{ll} x_r - x_r & \leq 0 \\ x_r - x_t & \leq 0 \\ x_s & \leq \frac{1}{4} \\ x_s - \frac{1}{2}x_t & \leq 0 \\ -\frac{1}{2}x_s + x_t & \leq \frac{1}{2} \end{array}$$

Fig. 1. A concurrent probabilistic system together with the linear optimization arising from the formula $\phi = [a\ \forall\mathcal{U}\ b]_{>\frac{1}{4}}$ without fairness.

the values of a function $f_A : \{1, \ldots 2^m\} \times \{1, \ldots 2^m\} \to D$, where $f_A(i,j) = a_{ij}$, mapping the position indices i, j to the matrix element value a_{ij}. Using the standard encoding $c : \{0,1\}^m \to \{1, \ldots 2^m\}$ of Boolean sequences of length m into the integers, this function may be interpreted as a D–valued Boolean function $f : \{0,1\}^{2m} \to D$ where $f(x,y) = f_A(c(x), c(y))$ for $x = (x_1 \ldots x_m)$ and $y = (y_1 \ldots y_m)$. Furthermore, we require the variables of f to alternate, thus using the MTBDD obtained from $f(x_1, y_1, x_2, y_2, \ldots, x_m, y_m)$ instead of the MTBDD for $f(x_1 \ldots x_m, y_1 \ldots y_m)$. This convention imposes a recursive structure on the matrix from which efficient recursive algorithms for all standard matrix operations are derived [12].

Matrices can be compactly represented by MTBDDs if they have submatrices that are identical; since we are dealing with very sparse probability matrices this applies to our case. Though in the worst case exponential, compared to sparse matrix representation MTBDDs provide a uniform $\log_2(N)$ access time (where N is the number of non-zero elements of the matrix) and combine efficiently with BDDs. Also, we expect MTBDDs to yield a practical advantage over sparse matrices depending on the degree of regularity of the original matrix. Alternative data structures, for example FEVBDDs [34] and DNBDDs [30, 31], are less well developed; in addition, it is not clear how to adapt FEVBDDs to floating point calculations.

We now describe the representation of concurrent probabilistic transition systems with MTBDDs, which generalise the labelled Markov chains of [4]. Consider a probabilistic transition system with n states that enable at most l probabilistic transitions each. Suppose also that n and l are powers of 2. We can represent the probabilistic transition system as a nl by n matrix T. Let m denote nl. Each row of T represents a single probabilistic transition, where the element in position $(i.k, j)$ represents the probability of reaching state j from state i in the k^{th} transition that leaves from i. If from a state i there are less than l transitions enabled, then we have several options, including repeating the last row (adopted here) or padding with all-zero rows (in which case care must be taken so that the zeros do not interfere with the calculations).

5 PBTL Model Checking with MTBDDs and Simplex

We now outline the symbolic model checking for PBTL formulas based on MTBDDs. We assume that the concurrent probabilistic transition system has been represented as an MTBDD as described above. Given a PBTL formula the algorithm proceeds by inductively converting the formulas to their BDD representation. With the exception of "until" formulas (and also "next state" not considered here) this is similar to the algorithm presented in [4].

For the PBTL formulas such as $[\psi\ \forall\mathcal{U}\ \psi]_{\sqsupseteq\lambda}$ we implement two methods: the iteration as described in Section 3 (a straightforward limiting process that computes the probability

to within a specified ϵ, omitted from this discussion) and by reduction to a linear programming problem which we solve by simplex. We now describe how to generate the linear programming problem.

5.1 Generation of the Linear Programming Problem

First we need to compute the set of states from which under any scheduling there is a non-zero probability of satisfying $\phi \; \mathcal{U} \; \psi$. We start by defining two Boolean column vectors $\mathbf{c\_b}_\phi$ and $\mathbf{c\_b}_\psi$ that represent the states satisfying ϕ and ψ, respectively. These correspond to the sets S_ϕ, S_ψ of Section 3. Observe that these vectors are independent of the bits representing transitions, and thus their BDD representations as n vectors and m vectors are the same. In the rest of the section we use $\mathbf{c\_b}_\phi$ and $\mathbf{c\_b}_\psi$ interchangeably as vectors and as functions depending on what is more convenient.

We define the set $REACH(\mathbf{c\_b}_\phi, \mathbf{c\_b}_\psi, T)$ as the minimum fixpoint $\mu Z[R]$ of the relation R defined by

$$R = \lambda x.(\mathbf{c\_b}_\psi[x] \vee (\mathbf{c\_b}_\phi[x] \wedge \forall_k \exists_y (Z[y] \wedge is_p ositive(T[x.k, y])))).$$

We anticipate that iterative squaring can be incorporated, although we have not considered this question as yet.

Now we let $\mathbf{c\_b}_{>\mathbf{0}}$ denote $REACH(\mathbf{c\_b}_\phi, \mathbf{c\_b}_\psi, T)$, the set of states that satisfy $[\phi \; \forall\mathcal{U} \; \psi]_{>0}$ ($S^{>0}$ of Section 3) and define

$$\mathbf{c\_b}_? = \lambda x.(\mathbf{c\_b}_{>\mathbf{0}}[x] \wedge \neg \mathbf{c\_b}_\psi[x])$$

which represents the set of states where the probability of satisfying $\phi \; \mathcal{U} \; \psi$ is still unknown ($S^?$ in Section 3). We also want to identify all the transitions that leave from states of $\mathbf{c\_b}_?$. To this purpose we define the Boolean vector

$$\mathbf{c\_t}_? = \lambda xk.(\mathbf{c\_b}_?[x]).$$

Observe that the rules of extending the domain of a function without changing the related BDD representation are still valid. All our arguments are up to such extensions. Indeed now we need to define a Boolean matrix that identifies the relevant elements of the transition matrix T:

$$B = \lambda xky.\mathbf{c\_t}_?[x.k] \wedge \mathbf{c\_b}_?[y],$$

which is essentially $\mathbf{c\_t}_? \cdot \mathbf{c\_b}_?{}^t$, where $(\cdot)^t$ denotes transpose. Taking care to match the sizes of T and B, define

$$T' = APPLY(T, B, *)$$

to be the pointwise multiplication between T and B [3], which represents the set of constraints for the linear programming problem to be solved. Let

$$I = \lambda xky.(x = y)$$

denote the identity matrix on the states, irrespective of the transition enabled. Then we define

$$A = APPLY(I, T', -)$$

A is the matrix part of the tableau. In order to use the the simplex algorithm on A we need to add slack variables: increase (double) the size of A to an m by $2m$ matrix by

[3] $APPLY(M_1, M_2, \oplus)$, where $\oplus$ is a binary operation on $[0, 1]$, is an operator which applies $\oplus$ pointwise to the matrices M_1 and M_2 in MTBDD form.

appending the identity matrix $I_{m,m}$ to the right-hand side of A. It is assumed that the unused columns are identically 0.

Next we compute a column vector called **c_q** (represented on the right- or left-hand side of the matrix in the usual tableau presentation of simplex, see e.g. [23]). Let

$$\mathbf{c\_q} = T\mathbf{c\_b}_{\psi}^{t}$$

express the probability that each transition reaches a state satisfying ψ. Recall that each line of T denotes one transition and that $\mathbf{c\_b}_{\psi}$ denotes the states that satisfy ψ.

The cost function is defined to be $\mathbf{r\_cost} = -\mathbf{c\_b}_{?}{}^{t}$, a row vector. The problem that we need to solve is to *maximize* $\mathbf{r\_cost} \cdot x$ under the constraint $Ax = \mathbf{c\_q}$. The vector x will denote the probabilities with which each state satisfies $\phi \,\mathcal{U}\, \psi$ (of course, only the x_i's that represent states).

We start by defining the base with all the slack variables, i.e., a row vector **base** with m 0's followed by m 1's. Then we implement the simplex algorithm with Bland's rule based on [23].

5.2 Implementing Simplex in MTBDDs

One iteration of the simplex algorithm is as follows.

First we find the column entering the basis – this is the column at the position of the first negative element of the cost vector. To do this we build a Boolean vector (called **c_enter**) whose only true element (1 element) is the smallest i such that $\mathbf{r\_cost}[i] < 0$. To do this, first we compute the pointwise predicate that tests for negativity, and find the position of the first negative element through a BDD operation.

Next we extract the actual column that enters the basis by computing $\mathbf{c\_enter\_d} = A \cdot \mathbf{c\_enter}$. Then we compute the pointwise division operation between **c_q** and **c_enter_d** where division by 0 gives ∞ and negative results are treated as ∞ as well. Finally, we compute the minimum index minimum element of the result of the division, leading to a vector **r_exit** with a structure similar to **c_enter** (**c_enter** has $2m$ elements; **r_exit** has m elements).

Then we extract the vector **r_exit_d** of the row identified by **r_exit** by computing $\mathbf{r\_exit\_d} = \mathbf{r\_exit} \cdot A$. Finally, we compute the vector that identifies the column that leaves the base as $\mathbf{r\_leave\_base} = \mathbf{base} \wedge \mathbf{r\_exit\_d}$ (a value different from 0 is regarded as true; note that only one element of **base** is different from 0 in **r_exit_d**), and we compute the new base as (**r_leave_base** *xor* **base**) $\vee$ **c_enter**, where all operations are taken pointwise. Note that only one element of **r_leave_base** is not 0.

It remains to change the basis: we know from **c_enter** what column will enter the base and we know from **r_exit** what row will represent the entering column. Our objective is to update the entering column so that all elements except for the element in the position identified by **r_exit** are 0 and the element identified by **r_exit** is 1. Furthermore, we need to update the cost vector so that it has a 0 in the position identified by **c_enter**. This is done through the following steps:

1. Calculate new matrix A of the tableau by subtracting from each row of A the row **r_exit_d** multiplied by an appropriate factor, then normalizing **r_exit_d** so that it has 1 in the position identified by **c_enter**.
2. Compute new cost **r_cost** by subtracting from it the row **r_exit_d** multiplied by an appropriate factor so that it has a 0 in the position identified by **c_enter**.
3. Compute new **c_q** similarly to **r_cost**.

The iteration process terminates as soon as optimal vector is found, i.e. the new **c_enter** is zero.

5.3 Computing Validity of "until" Formulas

Finally, given the solution vector x of probabilities we need to compute the validity of the initial formula $[\phi\ \forall\mathcal{U}\ \psi]_{\sqsupseteq\lambda}$. Taking appropriate care of the slack variables, this can be achieved by computing pointwise the operation $x[i] \geq \lambda$, which is an ordinary BDD operation.

The case of formulas of the form $[\phi\ \exists\mathcal{U}\ \psi]_{\sqsupseteq\lambda}$ can be dealt with by transposing the matrix A [4], the extra column and the cost function, then using the same simplex algorithm for thus reformulated problem which becomes: maximize $\mathbf{c\_q}^t \cdot y$ subject to the constraints $A^t y = \mathbf{r\_cost}^t$ (where A^t takes appropriate care of slack variables).

Example: For the PBTL-structure $\mathcal{S}$ of Figure 1 and $\phi = [a\ \exists\mathcal{U}\ b]_{>\frac{1}{4}}$. The linear programming problem that ϕ gives rise to is: minimize

$$x_r + x_s + x_t$$

subject to the constraints:

$$\begin{array}{ll} x_r - x_r & \geq 0 \quad (y_0) \\ x_r - x_t & \geq 0 \quad (y_1) \\ x_s & \geq \frac{1}{4} \quad (y_2) \\ x_s - \frac{1}{2}x_t & \geq 0 \quad (y_3) \\ -\frac{1}{2}x_s + x_t & \geq \frac{1}{2} \quad (y_4) \end{array}$$

and yields $x_r = \frac{2}{3}$, $x_s = \frac{1}{3}$ and $x_t = \frac{2}{3}$ as optimal. When reformulated to the dual problem, this reduces to: maximize

$$\frac{1}{4}y_2 + \frac{1}{2}y_4$$

subject to the constraints:

$$\begin{array}{ll} y_1 & \leq 1 \\ y_2 + y_3 - \frac{1}{2}y_4 & \leq 1 \\ -y_1 - \frac{1}{2}y_3 + y_4 & \leq 1 \end{array}$$

which can be solved by the same method as before. The values for x_r, x_s and x_t can be read off the optimal tableau.

The calculations needed for satisfaction with respect to fairness constraints follow by a similar method through the reduction to the equivalent optimization problem described in Section 3.2.

6 Experimental Results

We have implemented an experimental symbolic model checking tool in terms of the CUDD package [32]. This package provides support for BDDs and MTBDDs, together with the matrix multiplication as given in [1] and the times-plus method proposed by E. Clarke. The CUDD package uses its own memory management, and hence its memory requirement is relatively high (typically over 50MB for medium size examples). The running times, on the other hand, are fast (i.e. a fraction of a second for small examples run on an Ultra 10).

Our tool accepts as input a real-valued, sparse matrix description of a concurrent probabilistic system[5], the type of satisfaction required, and a PBTL formula ϕ, and outputs the

[4] Technically speaking, since A already contains slacks, we transpose its left-hand side half only, excluding the slack variables, and add new slacks.

[5] We defer the implementation of the system description language till later, in the meantime using the TIPPtool [21] to construct case studies.

states that are satisfied by ϕ. Our implementation relies on the matrix-by-vector multiplication obtained from the matrix-by-matrix multiplication supplied with the CUDD package (we use the above mentioned times-plus algorithm). We work with double-presision floating point arithmetic, which is standard for numerical calculations. Due to the approximate nature of probabilities we do not expect the need for infinite precision to arise. However, we anticipate that our algorithms can be modified to cater for rationals represented by integer fractions that would guarantee better precision, potentially at a higher performance cost.

The model checking of unbounded "until" properties has been implemented by two methods. The first is based on the Jacobi method for iterative refinement, which has the advantages that it relies only on the matrix-by-vector multiplication, does not change the matrix representing the system, is numerically stable (has no build up of round off errors) and admits improvements to numerical precision. However, it can be inefficient in cases of slow convergence (we have encountered examples requiring over 10,000 iterations). Our alternative, the simplex method, is numerically not as stable as it involves subtraction, and also destroys the original matrix. However, there are cases which require fewer iterations of simplex than the corresponding number of iterations.

We have tested the simplex component of our tool on several small and medium-size examples (such as the example in Figure 1, the randomized dining philosophers and the mutual exclusion [25]), and have verified all the liveness properties stated in [25]. The verification of these properties involves conventional BDD reachability analysis (cf the set $REACH(.,.,T)$ of Subsection 5.1) and Boolean operations, together with one linear optimization for each nested "until" property. It turns out that many liveness properties reduce to reachability analysis followed by a linear optimization that can be solved by simplex or iteration in 0 or 1 iterations.

Because the above problems are not a good test of the performance of simplex, we have set up an experiment based on the models as above, but using single universal "until" properties chosen so that both methods require a non-trivial number of iterations. For both simplex and the iterative refinement we compare the performance of our algorithms with equivalent MATLAB implementations using sparse matrices (note that we implement the basic tableaux simplex as in [23] in MATLAB). The results are very encouraging, see the summary below which shows the time in seconds on an Ultra 10 with 380MB and, in square brackets, the number of iterations required.

Example from Figure 1:

$[a\ \forall\mathcal{U}\ b]_{>\frac{1}{4}}$	Iterative		Simplex	
	Verif	Matlab	Verif	Matlab
$\models$	0.01 [4]	0.02 [4]	0.01 [2]	0.13 [2]
$\models_{fair}$	0.01 [4]	0.02 [4]	0.01 [3]	0.01 [3]

Randomized Dining Philosophers [25]:
3 Philosophers, 770 states and 2845 non-zeros in transition matrix.

Property	Iterative		Simplex	
	Verif	Matlab	Verif	Matlab
1	0.19 [19]	1.19 [19]	0.49 [59]	7.64 [59]
2	0.30 [19]	1.02 [19]	1.26 [154]	21.25 [154]
3	1.59 [44]	0.82 [44]	45.49 [360]	439 [360]
4	1.71 [49]	0.83 [49]	220 [404]	1092 [417]

4 Philosophers, 7070 states and 33405 non-zeros in transition matrix.

Property	Iterative		Simplex	
	Verif	Matlab	Verif	Matlab
1	0.28 [14]	102.42 [14]	0.94 [59]	58.23 [59]
2	1.79 [22]	108.60 [22]	20.44 [387]	833 [387]
3	1.73 [24]	97.41 [24]	25.44 [435]	1151 [435]

Randomized Mutual Exclusion [25]: 3 processes, 2680 states and 9646 non-zeros in transition matrix.

Property	Iterative		Simplex	
	Verif	Matlab	Verif	Matlab
1	3.73 [58]	15.08 [58]	8.19 [438]	261 [441]
2	4.78 [58]	13.21 [58]	17.27 [669]	722 [672]

Shortly we intend to investigate efficiency improvements such as SOR for the iterative method and the revised simplex method [27].

7 Conclusion

We have proposed MTBDD-based symbolic model checking for the probabilistic branching-time temporal logic PBTL of [7, 5] including fairness, and implemented an experimental tool using the CUDD package, with the view to investigate the feasibility of an integrated tool combining BDD-based model checking with probabilistic and performance analysis. Our experiments so far point to potential practical advantage in using MTBDDs as a representation for very sparse probability matrices. In particular, the MTBDD implementation consistently outperforms the equivalent MATLAB implementation of the basic simplex, with the iterative refinement being faster on the majority of the examples tested. This is because the model checking is a combination of BDD and MTBDD operations, and hence the gain is through the compactness of the representation and efficiency of the BDD operations.

It should be noted that, similarly to the Gaussian elimination described in [1], the structure of the simplex algorithm does not match the recursive structure of the MTBDD well (it works on a column-by-column basis, not subblock-by-subblock). In addition, it does not preserve the sparsity of the matrix (the matrix fills up as the computation progresses, which is evident in the time per iteration increasing). Nevertheless, for the very sparse matrices it performs unexpectedly well.

There are many efficiency improvements that are worth investigating, to mention heuristics concerning the variable ordering, caching policy, the use of iterative squaring, and further known improvements to the iteration and the simplex method, see e.g. [27]. Furthermore, we would like to implement an algorithm for computing stationary state probabilities for labelled Markov chains. A generalisation of the model checking procedure to incorporate alternative fairness notions, see e.g. [5, 14], is relatively straightforward.

Finally, we believe that the MTBDD version of the simplex algorithm that we have implemented is of independent interest. It would be interesting to investigate recursive analogues to linear algebra operations, and whether they can be efficiently coded with MTBDDs.

References

1. R. I. Bahar, E. A. Frohm, C. M. Gaona, G. D. Hachtel, E. Macii, A. Pardo, and F. Somenzi. Algebraic decision diagrams and their applications. *Journal of Formal Methods in Systems Design*, 10(2/3):171–206, 1997.
2. C. Baier. On algorithmic verification methods for probabilistic systems. Habilitation thesis, submitted, 1998.

3. C. Baier, E. Clarke, and V. Hartonas-Garmhausen. On the semantic foundations of Probabilistic VERUS. In *Proceedings, Probabilistic Methods in Verification (PROBMIV'98)*, volume 21 of *Electronic Notes in Theoretical Computer Science*. Elsevier, 1999. To appear.
4. C. Baier, E. Clarke, V. Hartonas-Garmhausen, M. Kwiatkowska, and M. Ryan. Symbolic model checking for probabilistic processes. In *Proceedings, 24th ICALP*, volume 1256 of *Lecture Notes in Computer Science*, pages 430–440. Springer-Verlag, 1997.
5. C. Baier and M. Kwiatkowska. Model checking for a probabilistic branching time logic with fairness. *Distributed Computing*, 11:125–155, 1998.
6. J. Bengtsson, K. G. Larsen, F. Larsson, P. Pettersson, W. Yi, and C. Weise. New generation of uppaal. In *Proceedings of the International Workshop on Software Tools for Technology Transfer*, Aalborg, Denmark, July 1998.
7. A. Bianco and L. de Alfaro. Model checking of probabilistic and nondeterministic systems. In *Proceedings, FST&TCS*, volume 1026 of *Lecture Notes in Computer Science*, pages 499–513. Springer-Verlag, 1995.
8. M. Bozga, C. Daws, O. Maler, A. Olivero, S. Tripakis, and S. Yovine. Kronos: a model-checking tool for real-time systems. In *Proc. of the 10th Conference on Computer-Aided Verification*, Vancouver, Canada, 28 June - 2 July 1998. Springer Verlag.
9. R. Bryant. Graph-based algorithms for Boolean function manipulation. *IEEE Transactions on Computers*, C-35(8):677–691, 1986.
10. J. R. Burch, E. M. Clarke, K. L. McMillan, D. L. Dill, and J. Hwang. Symbolic model checking: 10^{20} states and beyond. Proceedings of the Fifth Annual Symposium on Logic in Computer Science, June 1990.
11. E. Clarke, E. Emerson, and A. Sistla. Automatic verification of finite state concurrent systems using temporal logic specifications: A pratical approach. In *Proceedings, 10th Annual Symp. on Principles of Programming Languages*. ACM Press, 1983.
12. E. Clarke, M. Fujita, P. McGeer, J.Yang, and X. Zhao. Multi-Terminal Binary Decision Diagrams: An Efficient Data Structure for Matrix Representation. In *IWLS'93: International Workshop on Logic Synthesis, Tahoe City*, May 1993.
13. C. Courcoubetis and M. Yannakakis. The complexity of probabilistic verification. *Journal of the ACM*, 42(4):857–907, 1995.
14. L. de Alfaro. Stochastic transition systems. In *Proc. CONCUR'98: Concurrency Theory*, Lecture Notes in Computer Science. Springer-Verlag, 1998.
15. L. de Alfaro. From fairness to chance. In *Proceedings, Probabilistic Methods in Verification (PROBMIV'98)*, volume 21 of *Electronic Notes in Theoretical Computer Science*. Elsevier, 1999. To appear.
16. G. D. Hachtel, E. Macii, A. Pardo, and F. Somenzi. Markovian analysis of large finite state machines. *IEEE CAD*, 15(12):1479–1493, Dec. 1996.
17. H. Hansson. *Time and Probability in Formal Design of Distributed Systems*. Elsevier, 1994.
18. H. Hansson and B. Jonsson. A logic for reasoning about time and probability. *Formal Aspects of Computing*, 6:512–535, 1994.
19. S. Hart, M. Sharir, and A. Pnueli. Termination of probabilistic concurrent programs. *ACM Transactions on Programming Languages and Systems*, 5:356–380, 1983.
20. V. Hartonas-Garmhausen. *Probabilistic Symbolic Model Checking with Engineering Models and Applications*. PhD thesis, Carnegie Mellon University, 1998.
21. H. Hermanns, V. Mertsiotakis, and M. Rettelbach. A construction and analysis tool based on the Stochastic Process Algebra TIPP. In *Proceedings, TACAS Workshop*, volume 1055 of *Lecture Notes in Computer Science*. Springer-Verlag, 1996.
22. N. Higham. *Accuracy and Stability of Numerical Algorithms*. SIAM, 1996.
23. H. Karloff. *Linear Programming*. Birkhauser, 1991.
24. K. McMillan. *Symbolic Model Checking*. Kluwer Academic Publishers, 1993.
25. A. Pnueli and L. Zuck. Verification of multiprocess probabilistic protocols. *Distributed Computing*, 1:53–72, 1986.
26. A. Pnueli and L. Zuck. Probabilistic verification. *Information and Computation*, 103:1–29, 1993.
27. A. Schrijver. *Theory of Linear and Integer Programming*. J. Wiley and Sons, 1986.
28. R. Segala. *Modelling and Verification of Randomized Distributed Real-Time Systems*. PhD thesis, MIT, 1995.
29. R. Segala and N. Lynch. Probabilistic simulations for probabilistic processes. In *Proceedings, CONCUR*, volume 836 of *Lecture Notes in Computer Science*, pages 481–496. Springer-Verlag, 1994.

30. M. Siegle. Compact representation of large performability models based on extended BDDs. In *Fourth International Workshop on Performability Modeling of Computer and Communication Systems (PMCCS4)*, pages 77–80, Williamsburg, VA, September 1998.
31. M. Siegle. Technique and tool for symbolic representation and manipulation of stochastic transition systems. TR IMMD 7 2/98, Universität Erlangen-Nürnberg, March 1998. http://www7.informatik.uni-erlangen.de/~siegle/own.h tml.
32. F. Somenzi. CUDD: CU decision diagram package. Public software, Colorado University, Boulder, 1997.
33. W. Stewart. *Introduction to the Numerical Solution of Markov Chains*. Princeton University Press, 1994.
34. P. Tafertshofer and M. Pedram. Factored edge-valued binary-decision diagrams. *Formal Methods in Systems Design*, 10:137–164, 1997.
35. M. Vardi. Automatic verification of probabilistic concurrent finite-state programs. In *Proceedings, FOCS'85*, pages 327–338. IEEE Press, 1987.