Least Upper Bounds for Probability Measures and their Applications to Abstractions

Rohit Chadha[1⋆], Mahesh Viswanathan[1⋆⋆], and Ramesh Viswanathan[2]

[1] University of Illinois at Urbana-Champaign
{rch,vmahesh}@uiuc.edu
[2] Bell Laboratories
rv@research.bell-labs.com

Abstract. Abstraction is a key technique to combat the state space explosion problem in model checking probabilistic systems. In this paper we present new ways to abstract Discrete Time Markov Chains (DTMCs), Markov Decision Processes (MDPs), and Continuous Time Markov Chains (CTMCs). The main advantage of our abstractions is that they result in abstract models that are *purely probabilistic*, which maybe more amenable to automatic analysis than models with both nondeterministic and probabilistic steps that typically arise from previously known abstraction techniques. A key technical tool, developed in this paper, is the construction of least upper bounds for any collection of probability measures. This upper bound construction may be of independent interest that could be useful in the abstract interpretation and static analysis of probabilistic programs.

1 Introduction

Abstraction is an important technique to combat *state space explosion*, wherein a smaller, abstract model that conservatively approximates the behaviors of the original (concrete) system is verified/model checked. Typically abstractions are constructed on the basis of an equivalence relation (of finite index) on the set of (concrete) states of the system. The abstract model has as states the equivalence classes (*i.e.*, it collapses all equivalent states into one), and each abstract state has transitions corresponding to the transitions of each of the concrete states in the equivalence class. Thus, the abstract model has both nondeterministic and (if the concrete system is probabilistic) probabilistic behavior.

In this paper, we present new methods to abstract probabilistic systems modeled by Discrete Time Markov Chains (DTMC), Markov Decision Processes (MDP), and Continuous Time Markov Chains (CTMC). The main feature of our constructions is that the resulting abstract models are *purely probabilistic* in that they do not have any nondeterministic choices. Since analyzing models that have both nondeterministic and probabilistic behavior is more challenging than

⋆ Supported partially by NSF CCF 04-29639
⋆⋆ Supported partially by NSF CCF 04-48178

analyzing models that are purely probabilistic, we believe that this may make our abstractions more amenable to automated analysis; the comparative tractability of model-checking systems without non-determinism is further detailed later in this section.

Starting from the work of Saheb-Djahromi [19], and further developed by Jones [11], orders on measures on special spaces (Borel sets generated by Scott open sets of a cpo) have been used in defining the semantics of probabilistic programs. Ordering between probability measures also play a central role in defining the notion of simulation for probabilistic systems. For a probabilistic model, a transition can be viewed as specifying a probability measure on successor states. One transition then simulates another if the probability measures they specify are related by the ordering on measures. In this manner, simulation and bisimulation relations were first defined for DTMCs and MDPs [12], and subsequently extended to CTMCs [3]. Therefore, in all these settings, a set of transitions is abstracted by a transition if it is an upper bound for the probability measures specified by the set of transitions being abstracted.

The key technical tool that we develop in this paper is a new construction of least upper bounds for arbitrary sets of probability measures. We show that in general, measures (even over simple finite spaces) do not have least upper bounds. We therefore look for a class of measurable spaces for which the existence of least upper bounds is guaranteed for arbitrary sets of measures. Since the ordering relation on measures is induced from the underlying partial order on the space over which the measures are considered, we identify conditions on the underlying partial order that are sufficient to prove the existence of least upper bounds — intuitively, these conditions correspond to requiring the Hasse diagram of the partial order to have a "tree-like" structure. Furthermore, we show that these conditions provide an exact characterization of the measurable spaces of our interest — for any partial order not satisfying these conditions, we can construct two measures that do not have a least upper bound. Finally, for this class of tree-like partial orders, we provide a natural construction that is proven to yield a well-defined measure that is a least upper bound.

These upper bound constructions are used to define abstractions as follows. As before, the abstract model is defined using an equivalence relation on the concrete states. The abstract states form a tree-like partial order with the minimal elements consisting of the equivalence classes of the given relation. The transition out of an abstract state is constructed as the least upper bound of the transitions from each of the concrete states it "abstracts". Since each upper bound is a single measure yielding a single outgoing transition, the resulting abstract model does not have any nondeterminism. This intuitive idea is presented and proved formally in the context of DTMCs, MDPs and CTMCs.

A few salient features of our abstract models bear highlighting. First, the fact that least upper bounds are used in the construction implies that for a particular equivalence relation on concrete states and partial order on the abstract states, the abstract model constructed is finer than (*i.e.*, can be simulated by) any purely probabilistic models that can serve as an abstraction. Thus, for veri-

fication purposes, our model is the most precise purely probabilistic abstraction on a chosen state space. Second, the set of abstract states is not completely determined by the equivalence classes of the relation on concrete states; there is freedom in the choice of states that are above the equivalence classes in the partial order. However, for any such choice that respects the "tree-like" requirement on the underlying partial order, the resulting model will be exponentially smaller than the existing proposals of [8, 13]. Furthermore, we show that there are instances where we can get more precise results than the abstraction schemes of [8, 13] while using significantly fewer states (see Example 4). Third, the abstract models we construct are purely probabilistic which makes model checking easier. Additionally, these abstractions can potentially be verified using statistical techniques which do not work when there is nondeterminism [24, 23, 21]. Finally, CTMC models with nondeterminism, called CTMDP, are known to be difficult to analyze [2]. Specifically, the measure of time-bounded reachability can only be computed approximately through an iterative process; therefore, there is only an approximate algorithm for model-checking CTMDPs against CSL. On the other hand, there is a theoretically exact solution to the corresponding model-checking problem for CTMCs by reduction to the first order theory of reals [1].

Related Work. Abstractions have been extensively studied in the context of probabilistic systems. General issues and definitions of good abstractions are presented in [12, 9, 10, 17]. Specific proposals for families of abstract models include Markov Decision Processes [12, 5, 6], systems with interval ranges for transition probabilities [12, 17, 8, 13], abstract interpretations [16], 2-player stochastic games [14], and expectation transformers [15]. Recently, theorem-prover based algorithms for constructing abstractions of probabilistic systems based on predicates have been presented [22]. All the above proposals construct models that exhibit both nondeterministic and probabilistic behavior. The abstraction method presented in this paper construct purely probabilistic models.

2 Least Upper Bounds for Probability Measures

This section presents our construction of least upper bounds for probability measures. Section 2.1 recalls the relevant definitions and results from measure theory. Section 2.2 presents the ordering relation on measures. Finally, Section 2.3 presents the least upper bound construction on measures. Due to space considerations, many of the proofs are deferred to [4] for the interested reader.

2.1 Measures

A **measurable space** (X, Σ) is a set X together with a family of subsets, Σ, of X, called a **σ-field** or **σ-algebra**, that contains $\emptyset$ and is closed under complementation and countable union. The members of a σ-field are called the **measurable subsets** of X. Examples of σ-fields are $\{\emptyset, X\}$ and $\mathcal{P}(X)$ (the powerset of X). We will sometimes abuse notation and refer to the measurable

space (X, Σ) by X or by Σ, when the σ-field or set, is clear from the context. The intersection of an arbitrary collection of σ-fields on a set X is again a σ-field, and so given any $B \subseteq \mathcal{P}(X)$ there is a least σ-field containing B, which is called the σ-field **generated** by B.

A **positive measure** μ on a measurable space (X, Σ) is a function from Σ to $[0, \infty]$ such that μ is **countably additive**, *i.e.*, if $\{A_i \mid i \in I\}$ is a countable family of pairwise disjoint measurable sets then $\mu(\bigcup_{i \in I} A_i) = \sum_{i \in I} \mu(A_i)$. In particular, if $I = \emptyset$, we have $\mu(\emptyset) = 0$. A measurable space equipped with a measure is called a **measure space**. The **total weight** of a measure μ on measurable space X is $\mu(X)$. A **probability measure** is a positive measure of total weight 1. We denote the collection of all probability measures on X by $\mathcal{M}_{=1}(X)$.

2.2 A Partial Order on Measures

In order to define an ordering on probability measures we need to consider measurable spaces that are equipped with an ordering relation. An *ordered measurable space* $(X, \Sigma, \sqsubseteq)$ is a set X equipped with a σ-field Σ and a preorder on X [3] $\sqsubseteq$ such that (X, Σ) is a measurable space. A (probability) measure on $(X, \Sigma, \sqsubseteq)$ is a (probability) measure on (X, Σ). Finally, recall that a set $U \subseteq X$ is *upward closed* if for every $x \in U$ and $y \in X$ with $x \sqsubseteq y$ we have that $y \in U$. The ordering relation on the underlying set is lifted to an ordering relation on probability measures as follows.

Definition 1. *Let $\mathcal{X} = (X, \Sigma, \sqsubseteq)$ be an ordered measurable space. For any probability measures μ, ν on $\mathcal{X}$, define $\mu \leq \nu$ iff for every upward closed set $U \in \Sigma$, $\mu(U) \leq \nu(U)$.*

Our definition of the ordering relation is formulated so as to be applicable to any general measurable space. For probability distributions over finite spaces, it is equivalent to a definition of *lifting of preorders to probability measures* using *weight functions* as considered in [12] for defining simulations. Indeed, Definition 1 can be seen to be identical to the presentation of the simulation relation in [7, 20] where this equivalence has been observed as well.

Recall that a set $D \subseteq X$ is *downward closed* if for every $y \in D$ and $x \in X$ with $x \sqsubseteq y$ we have that $x \in D$. The ordering relation on probability measures can be dually cast in terms of downward closed sets which is useful in the proofs of our construction.

Proposition 1. *Let $\mathcal{X} = (X, \Sigma, \sqsubseteq)$ be an ordered measurable space. For any probability measures μ, ν on $\mathcal{X}$, we have that $\mu \leq \nu$ iff for every downward closed set $D \in \Sigma$, $\mu(D) \geq \nu(D)$.*

[3] Recall that preorder on a set X is a binary relation $\sqsubseteq \subseteq X \times X$ such that $\sqsubseteq$ is reflexive and transitive.

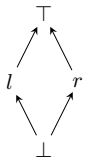

Fig. 1. Hasse Diagram of T. Arrows directed from smaller element to larger element.

In general, Definition 1 yields a preorder that is not necessarily a partial order. We identify a special but broad class of ordered measurable spaces for which the ordering relation is a partial order. The spaces we consider are those which are generated by some collection of downward closed sets.

Definition 2. *An ordered measurable space* $(X, \Sigma, \sqsubseteq)$ *is* order-respecting *if there exists* $\mathcal{D} \subseteq \mathcal{P}(X)$ *such that every* $D \in \mathcal{D}$ *is downward closed (with respect to* $\sqsubseteq$*) and* Σ *is generated by* $\mathcal{D}$*.*

Example 1. For any finite set A, the space $(\mathcal{P}(A), \mathcal{P}(\mathcal{P}(A)), \subseteq)$ is order-respecting since it is generated by all the downward closed sets of $(\mathcal{P}(A), \subseteq)$. One special case of such a space that we will make use of in our examples is where $T = \mathcal{P}(\{0, 1\})$ whose Hasse diagram is shown in Figure 1; we will denote the elements of T by $\bot = \emptyset$, $l = \{0\}$, $r = \{1\}$, and $\top = \{0, 1\}$. Then $\mathbb{T} = (T, \mathcal{P}(T), \subseteq)$ is an order-respecting measurable space. Finally, for any cpo $(X, \sqsubseteq)$, the Borel measurable space $(X, \mathcal{B}(X), \sqsubseteq)$ is order-respecting since every Scott-closed set is downward closed.

Theorem 1. *For any ordered measurable space* $\mathcal{X} = (X, \Sigma, \sqsubseteq)$*, the relation* $\leq$ *is a preorder on* $\mathcal{M}_{=1}(\mathcal{X})$*. The relation* $\leq$ *is additionally a partial order (anti-symmetric) if* $\mathcal{X}$ *is order-respecting.*

Example 2. Recall the space $\mathbb{T} = (T, \mathcal{P}(T), \subseteq)$ defined in Example 1. Consider the probability measure λ, where l has probability 1, and all the rest have probability 0. Similarly, τ is the measure where $\top$ has probability 1, and the rest 0, and in ρ, r gets probability 1, and the others 0. Now one can easily see that $\lambda \leq \tau$ and $\rho \leq \tau$. However $\lambda \not\leq \rho$ and $\rho \not\leq \lambda$.

2.3 Construction of Least Upper Bounds

Least upper bound constructions for elements in a partial order play a crucial role in defining the semantics of languages as well as in abstract interpretation. As we shall show later in this paper, least upper bounds of probabilistic measures can also be used to define abstract models of probabilistic systems. Unfortunately, however, probability measures over arbitrary measurable spaces do not necessarily have least upper bounds; this is demonstrated in the following example.

Example 3. Consider the space $\mathbb{T}$ defined in Example 1. Let μ be the probability measure that assigns probability $\frac{1}{2}$ to $\bot$ and l, and 0 to everything else. Let ν be such that it assigns $\frac{1}{2}$ to $\bot$ and r, 0 to everything else. The measure τ that assigns $\frac{1}{2}$ to $\top$ and $\bot$ is an upper bound of both μ and ν. In addition, ρ that assigns $\frac{1}{2}$ to l and r, and 0 to everything else, is also an upper bound. However τ and ρ are incomparable. Moreover, any lower bound of τ and ρ must assign a probability at least $\frac{1}{2}$ to $\bot$ and probability 0 to $\top$, and so cannot be an upper bound of μ and ν. Thus, μ and ν do not have a least upper bound.

We therefore identify a special class of ordered measurable spaces over which probability measures admit least upper bounds. Although our results apply to general measurable spaces, for ease of understanding, the main presentation here is restricted to finite spaces. For the rest of the section, fix an ordered measurable space $\mathcal{X} = (X, \mathcal{P}(X), \sqsubseteq)$, where $(X, \sqsubseteq)$ is a finite partial order. For any element $a \in X$, we use D_a to denote the downward-closed set $\{b \mid b \sqsubseteq a\}$. We begin by defining a *tree-like* partial order; intuitively, these are partial orders whose Hasse diagram resembles a tree (rooted at its greatest element).

Definition 3. *A partial order $(X, \sqsubseteq)$ is said to be* tree-like *if and only if (i) X has a greatest element $\top$, and (ii) for any two elements $a, b \in X$ if $D_a \cap D_b \neq \emptyset$ then either $D_a \subseteq D_b$ or $D_b \subseteq D_a$.*

We can show that over spaces whose underlying ordering is tree-like, any set of probability measures has a least upper bound. This construction is detailed in Theorem 2 and its proof below.

Theorem 2. *Let $\mathcal{X} = (X, \mathcal{P}(X), \sqsubseteq)$ be an ordered measurable space where $(X, \sqsubseteq)$ is tree-like. For any $\Gamma \subseteq \mathcal{M}_{=1}(\mathcal{X})$, there is a probability measure $\nabla(\Gamma)$ such that $\nabla(\Gamma)$ is the least upper bound of Γ.*

Proof. Recall that for a set $S \subseteq X$, its set of maximal elements maximal(S) is defined as $\{a \in S \mid \forall b \in S.\ a \sqsubseteq b \Rightarrow a = b\}$. For any downward closed set D, we have that $D = \cup_{a \in \text{maximal}(D)} D_a$. From condition (ii) of Definition 3, if a, b are two distinct maximal elements of a downward closed set D then $D_a \cap D_b = \emptyset$ and the sets comprising the union are pairwise disjoint. For any measure μ, we therefore have that $\mu(D) = \sum_{a \in \text{maximal}(D)} \mu(D_a)$ for any downward closed set D.

Define the function ν on downward closed subsets of X as follows. For a downward closed set of the form D_a, where $a \in X$, take $\nu(D_a) = \inf_{\mu \in \Gamma} \mu(D_a)$, and for any downward closed set D take $\nu(D) = \sum_{a \in \text{maximal}(D)} \nu(D_a)$. We will define the least upper bound measure $\nabla(\Gamma)$ by specifying its value pointwise on each element of X. Observe that for any $a \in X$, the set $D_a \setminus \{a\}$ is also downward closed. We therefore define $\nabla(\Gamma)(\{a\}) = \nu(D_a) - \nu(D_a \setminus \{a\})$, for any $a \in X$.

Observe that $\nu(D) \leq \inf_{\mu \in \Gamma} \mu(D)$. We therefore have that $\nabla(\Gamma)(\{a\}) \geq 0$. For any downward closed set D, we can see that $\nabla(\Gamma)(D) = \nu(D)$. Thus, $\nabla(\Gamma)(X) = \nabla(\Gamma)(D_\top) = \nu(D_\top) = \inf_{\mu \in \Gamma} \mu(D_\top) = 1$, and so $\nabla(\Gamma)$ is a probability measure on $\mathcal{X}$.

For any downward closed set D, we have that $\nabla(\Gamma)(D) = \nu(D)$ and $\nu(D) \leq \inf_{\mu \in \Gamma} \mu(D)$ which allows us to conclude that $\nabla(\Gamma)$ is an upper bound of Γ. Now consider any measure τ that is an upper bound of Γ. Then, $\tau(D) \leq \mu(D)$ for any measure $\mu \in \Gamma$ and all downward closed sets D. In particular, for any element $a \in X$, $\tau(D_a) \leq \inf_{\mu \in \Gamma} \mu(D_a) = \nu(D_a) = \nabla(\Gamma)(D_a)$. Thus, for any downward closed set D, we have that $\tau(D) = \sum_{a \in \text{maximal}(D)} \tau(D_a) \leq \sum_{a \in \text{maximal}(D)} \nabla(\Gamma)(D_a) = \nabla(\Gamma)(D)$. Hence, $\nabla(\Gamma) \leq \tau$, which concludes the proof. □

We conclude this section, by showing that if we consider any ordered measurable space that is not tree-like, there are measures that do not have least upper bounds. Thus, the tree-like condition is an *exact*(necessary and sufficient) characterization of spaces that admit least upper bounds of arbitrary sets of probability measures.

Theorem 3. *Let $\mathcal{X} = (X, \mathcal{P}(X), \sqsubseteq)$ be an ordered measurable space, where $(X, \sqsubseteq)$ is a partial order that is* not *tree-like. Then there are probability measures μ and ν such that μ and ν do not have a least upper bound.*

Proof. First consider the case when X does not have a greatest element. Then there are two maximal elements, say a and b. Let μ be the measure that assigns measure 1 to a and 0 to everything else, and let ν be the measure that assigns 1 to b and 0 to everything else. Clearly, μ and ν do not have an upper bound.

Next consider the case when X does have a greatest element $\top$; the proof in this case essentially follows from generalizing Example 3. If $\mathcal{X}$ is a space as in the theorem then since $(X, \sqsubseteq)$ is not tree-like, there are two elements $a, b \in X$ such that $D_a \cap D_b \neq \emptyset$, $D_a \setminus D_b \neq \emptyset$, and $D_b \setminus D_a \neq \emptyset$. Let $c \in D_a \cap D_b$. Consider the measure μ to be such that $\mu(\{c\}) = \frac{1}{2}$, $\mu(\{a\}) = \frac{1}{2}$, and is equal to 0 on all other elements. Define the measure ν to be such that $\nu(\{c\}) = \frac{1}{2}$, $\nu(\{b\}) = \frac{1}{2}$, and is equal to 0 on all other elments. As in Example 3, we can show that μ and ν have two incomparable minimal upper bounds. □

Remark 1. All the results presented in the section can be extended to ordered measure spaces $\mathcal{X} = (X, \mathcal{P}(X), \sqsubseteq)$ when X is a countable set; see [4].

3 Abstracting DTMCs and MDPs

In this section we outline how our upper bound construction can be used to abstract MDPs and DTMCs using DTMCs. We begin by recalling the definitions of these models along with the notion of simulation and logic preservation in Section 3.1, before presenting our proposal in Section 3.2.

3.1 Preliminaries

We recall 3-valued PCTL and its discrete time models. In 3-valued logic, a formula can evaluate to either *true* ($\top$), *false* ($\bot$), or *indefinite* (?); let $\mathbb{B}_3 =$

$\{\bot, ?, \top\}$. The formulas of PCTL are built up over a finite set of atomic propositions AP and are inductively defined as follows.

$$\varphi ::= \text{true} \mid a \mid \neg\varphi \mid \varphi \wedge \varphi \mid \mathcal{P}_{\bowtie p}(X\varphi) \mid \mathcal{P}_{\bowtie p}(\varphi \,\mathcal{U}\, \varphi)$$

where $a \in \text{AP}$, $\bowtie \in \{<, \leq, >, \geq\}$, and $p \in [0, 1]$.

The models of these formulas are interpreted over *Markov Decision Processes*, formally defined as follows. Let Q be a finite set of states and let $\mathcal{Q} = (Q, \mathcal{P}(Q))$ be a measure space. A Markov Decision Process (MDP) $\mathcal{M}$ is a tuple $(Q, \rightarrow, L)$, where $\rightarrow \subseteq Q \times \mathcal{M}_{=1}(\mathcal{Q})$ (non-empty and finite), and $L : (Q \times \text{AP}) \rightarrow \mathbb{B}_3$ is a labeling function that assigns a value in $\mathbb{B}_3$ to each atomic proposition in each state. We will say $q \rightarrow \mu$ to mean $(q, \mu) \in \rightarrow$. A *Discrete Time Markov Chain* (DTMC) is an MDP with the restriction that for each state q there is exactly one probability measure μ such that $q \rightarrow \mu$. The 3-valued semantics of PCTL associates a truth value in $\mathbb{B}_3$ for each formula φ in a state q of the MDP; we denote this by $[\![q, \varphi]\!]_{\mathcal{M}}$. We skip the formal semantics in the interests of space and the reader is referred to [8] [4].

Theorem 4 (Fecher-Leucker-Wolf [8]). *Given an MDP $\mathcal{M}$, and a PCTL formula φ, the value of $[\![q, \varphi]\!]_{\mathcal{M}}$ for each state q, can be computed in time polynomial in $|\mathcal{M}|$ and linear in $|\varphi|$, where $|\mathcal{M}|$ and $|\varphi|$ denote the sizes of $\mathcal{M}$ and φ, respectively.*

Simulation for MDPs, originally presented in [12] and adapted to the 3-valued semantics in [8], is defined as follows. A preorder $\sqsubseteq \subseteq Q \times Q$ is said to be a *simulation* iff whenever $q_1 \sqsubseteq q_2$ the following conditions hold.

- If $L(q_2, a) = \top$ or $\bot$ then $L(q_1, a) = L(q_2, a)$ for every proposition $a \in \text{AP}$,
- If $q_1 \rightarrow \mu_1$ then there exists μ_2 such that $q_2 \rightarrow \mu_2$ and $\mu_1 \leq \mu_2$, where μ_1 and μ_2 are viewed as probability measures over the ordered measurable space $(Q, \mathcal{P}(Q), \sqsubseteq)$.

We say $q_1 \preceq q_2$ iff there is a simulation $\sqsubseteq$ such that $q_1 \sqsubseteq q_2$. A state q_1 in an MDP $(Q_1, \rightarrow_1, L_1)$ is simulated by a state q_2 in MDP $(Q_2, \rightarrow_2, L_2)$ iff there is a simulation $\sqsubseteq$ on the direct sum of the two MDP's (defined in the natural way) such that $(q_1, 0) \sqsubseteq (q_2, 1)$.

Remark 2. The ordering on probability measures used in simulation definition presented in [12, 8] is defined using *weight functions*. However, the definition presented here is equivalent, as has been also observed in [7, 20].

Finally, there is a close correspondence between simulation and the satisfaction of PCTL formulas according to the 3-valued interpretation.

Theorem 5 (Fecher-Leucker-Wolf [8]). *Consider q, q' states of MDP $\mathcal{M}$ such that $q \preceq q'$. For any formula φ, if $[\![q', \varphi]\!]_{\mathcal{M}} \neq ?$ then $[\![q, \varphi]\!]_{\mathcal{M}} = [\![q', \varphi]\!]_{\mathcal{M}}$* [5].

[4] In [8] PCTL semantics for MDPs is not given; however, this is very similar to the semantics for AMCs which is given explicitly.

[5] This theorem is presented only for AMC. But its generalization to MDPs can be obtained from the main observations given in [8]. See [4].

3.2 Abstraction by DTMCs

Abstraction, followed by progressive refinement, is one way to construct a small model that either proves the correctness of the system or demonstrates its failure to do so. Typically, the abstract model is defined with the help of an equivalence relation on the states of the system. Informally, the construction proceeds as follows. For an MDP/DTMC $\mathcal{M} = (Q, \rightarrow, L)$ and equivalence relation $\equiv$ on Q, the abstraction is the MDP $\mathcal{A} = (Q_{\mathcal{A}}, \rightarrow_{\mathcal{A}}, L_{\mathcal{A}})$, where $Q_{\mathcal{A}} = \{[q]_{\equiv} \mid q \in Q\}$ is the set of equivalence classes of Q under $\equiv$, and $[q]_{\equiv}$ has a transition corresponding to each $q' \rightarrow \mu$ for $q' \in [q]_{\equiv}$.

However, as argued by Fecher-Leucker-Wolf [8], model checking $\mathcal{A}$ directly may not be feasible because it has large number of transitions and therefore a large size. It maybe beneficial to construct a further abstraction of $\mathcal{A}$ and analyze the further abstraction. In what follows, we have an MDP, which maybe obtained as outlined above, that we would like to (further) abstract; for the rest of this section let us fix this MDP to be $\mathcal{A} = (Q_{\mathcal{A}}, \rightarrow_{\mathcal{A}}, L_{\mathcal{A}})$. We will first present the Fecher-Leucker-Wolf proposal, then ours, and compare the approaches, discussing their relative merits.

Fecher et al. suggest that a set of transitions be approximated by *intervals* that bound the probability of transitioning from one state to the next, according to any of the non-deterministic choices present in $\mathcal{A}$. The resulting abstract model, which they call an *Abstract Markov Chain* (AMC) is formally defined as follows.

Definition 4. *The Abstract Markov Chain (AMC) associated with $\mathcal{A}$ is formally the tuple $\mathcal{M} = (Q_{\mathcal{M}}, \rightarrow_{\ell}, \rightarrow_{u}, L_{\mathcal{M}})$, where $Q_{\mathcal{M}} = Q_{\mathcal{A}}$ is the set of states, and $L_{\mathcal{M}} = L_{\mathcal{A}}$ is the labeling function on states. The lower bound transition $\rightarrow_{\ell}$ and upper bound transition $\rightarrow_{u}$ are both functions of type $Q_{\mathcal{M}} \rightarrow (Q_{\mathcal{M}} \rightarrow [0,1])$, and are defined as follows:*

$$q \rightarrow_{\ell} \mu \text{ iff } \forall q' \in Q_{\mathcal{M}}.\ \mu(q') = \min_{q \rightarrow_{\mathcal{A}} \nu} \nu(\{q'\})$$
$$q \rightarrow_{u} \mu \text{ iff } \forall q' \in Q_{\mathcal{M}}.\ \mu(q') = \max_{q \rightarrow_{\mathcal{A}} \nu} \nu(\{q'\})$$

Semantically, the AMC $\mathcal{M}$ is interpreted as an MDP having from each state q any transition $q \rightarrow \nu$, where ν is a probability measure that respects the bounds defined by $\rightarrow_{\ell}$ and $\rightarrow_{u}$. More precisely, if $q \rightarrow_{\ell} \mu_{\ell}$ and $q \rightarrow_{u} \mu_{u}$ then $\mu_{\ell} \leq \nu \leq \mu_{u}$, where $\leq$ is to be interpreted as pointwise ordering on functions.

Fecher et al. demonstrate that the AMC $\mathcal{M}$ (defined above) does indeed simulate $\mathcal{A}$, and using Theorem 5 the model checking results of $\mathcal{M}$ can be reflected to $\mathcal{A}$. The main advantage of $\mathcal{M}$ over $\mathcal{A}$ is that $\mathcal{M}$ can be model checked in time that is a polynomial in $2^{|Q_{\mathcal{M}}|} = 2^{|Q_{\mathcal{A}}|}$; model checking $\mathcal{A}$ may take time more than polynomial in $2^{|Q_{\mathcal{A}}|}$, depending on the number of transitions out of each state q.

We suggest using the upper bound construction, presented in Section 2.3, to construct *purely probabilistic* abstract models that do not have any nondeterminism. Let $(X, \sqsubseteq)$ be a tree-like partial order. Recall that the set of minimal

elements of X, denoted by minimal(X), is given by minimal$(X) = \{x \in X \mid \forall y \in X.\ y \sqsubseteq x \Rightarrow x = y\}$.

Definition 5. *Consider the MDP $\mathcal{A} = (Q_{\mathcal{A}}, \rightarrow_{\mathcal{A}}, L_{\mathcal{A}})$. Let $(Q, \sqsubseteq)$ be a tree-like partial order, such that* minimal$(Q) = Q_{\mathcal{A}}$. *Let $\mathcal{Q} = (Q, \mathcal{P}(Q), \sqsubseteq)$ be the ordered measurable space over Q. Define the DTMC $\mathcal{D} = (Q_{\mathcal{D}}, \rightarrow_{\mathcal{D}}, L_{\mathcal{D}})$, where*

- $Q_{\mathcal{D}} = Q$,
- *For $q \in Q_{\mathcal{D}}$, let $\Gamma_q = \{\mu \mid \exists q' \in Q_{\mathcal{A}}.\ q' \sqsubseteq q$ and $q' \rightarrow_{\mathcal{A}} \mu\}$. Now, $q \rightarrow_{\mathcal{D}} \nabla(\Gamma_q)$, and*
- *For $q \in Q_{\mathcal{D}}$ and $a \in$ AP, if for any $q_1, q_2 \in Q_{\mathcal{A}}$ with $q_1 \sqsubseteq q$ and $q_2 \sqsubseteq q$, we have $L(q_1, a) = L(q_2, a)$ then $L(q, a) = L(q_1, a)$. Otherwise $L(q, a) =?$*

Proposition 2. *The DTMC $\mathcal{D}$ simulates the MDP $\mathcal{A}$, where $\mathcal{A}$ and $\mathcal{D}$ are as given in Definition 5.*

Proof. Consider the relation $R_{\sqsubseteq}$ over the states of the disjoint union of $\mathcal{A}$ and $\mathcal{D}$, defined as $R_{\sqsubseteq} = \{((q,0),(q,0)) \mid q \in Q_{\mathcal{A}}\} \cup \{((q',1),(q'',1)) \mid q', q'' \in Q_{\mathcal{D}}, q' \sqsubseteq q''\} \cup \{((q,0),(q',1)) \mid q \in Q_{\mathcal{A}}, q' \in Q_{\mathcal{D}}, q \sqsubseteq q'\}$. ¿From the definition of $\mathcal{D}$, definition of simulation and the fact that ∇ is the least upper bound operator, it can be shown that $R_{\sqsubseteq}$ is a simulation. □

The minimality of our upper bound construction actually allows to conclude that $\mathcal{D}$ is as good as any DTMC abstraction can be on a given state space. This is stated precisely in the next proposition.

Proposition 3. *Let $\mathcal{A} = (Q_{\mathcal{A}}, \rightarrow_{\mathcal{A}}, L_{\mathcal{A}})$ be an MDP and $(Q, \sqsubseteq)$ be a tree-like partial order, such that* minimal$(Q) = Q_{\mathcal{A}}$. *Consider the DTMC $\mathcal{D} = (Q_{\mathcal{D}}, \rightarrow_{\mathcal{D}}, L_{\mathcal{D}})$, as given in Definition 5. If $\mathcal{D}' = (Q_{\mathcal{D}}, \rightarrow'_{\mathcal{D}}, L_{\mathcal{D}})$ is a DTMC such that the relation $R_{\sqsubseteq}$ defined in the proof of Proposition 2 is a simulation of $\mathcal{A}$ by $\mathcal{D}'$ then $\mathcal{D}'$ simulates $\mathcal{D}$ also.*

Comparison with Abstract Markov Chains. Observe that any tree-like partial order $(Q, \sqsubseteq)$ such that minimal$(Q) = Q_{\mathcal{A}}$ is of size at most $O(|Q_{\mathcal{A}}|)$; thus, in the worst case the time to model check $\mathcal{D}$ is exponentially smaller than the time to model check $\mathcal{M}$. Further, we have freedom to pick the partial order $(Q, \sqsubseteq)$. The following proposition says that adding more elements to the partial order on the abstract states does indeed result in a refinement.

Proposition 4. *Let $\mathcal{A} = (Q_{\mathcal{A}}, \rightarrow_{\mathcal{A}}, L_{\mathcal{A}})$ be an MDP and $(Q_1, \sqsubseteq_1)$ and $(Q_2, \sqsubseteq_2)$ be tree-like partial orders such that $Q_1 \subseteq Q_2$, $\sqsubseteq_2 \cap (Q_1 \times Q_1) = \sqsubseteq_1$, and $Q_{\mathcal{A}} =$* minimal$(Q_1) =$ minimal(Q_2). *Let $\mathcal{D}_1$ be a DTMC over $(Q_1, \sqsubseteq_1)$ and $\mathcal{D}_2$ a DTMC over $(Q_2, \sqsubseteq_2)$ as in Definition 5. Then, $\mathcal{D}_1$ simulates $\mathcal{D}_2$.*

Thus, one could potentially identify the appropriate tree-like partial order to be used for the abstract DTMC through a process of abstraction-refinement.

Finally, we can demonstrate that even though the DTMC $\mathcal{D}$ is exponentially more succinct than the AMC $\mathcal{M}$, there are examples where model checking $\mathcal{D}$ can give a more precise answer than $\mathcal{M}$.

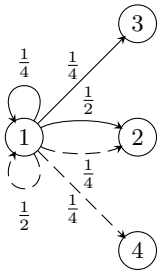

Fig. 2. Example MDP $\mathcal{A}$

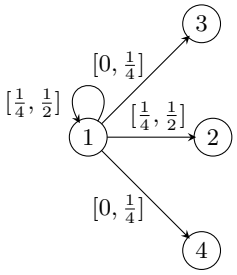

Fig. 3. AMC $\mathcal{M}$ corresponding to MDP $\mathcal{A}$

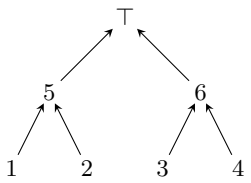

Fig. 4. Hasse diagram of partial order

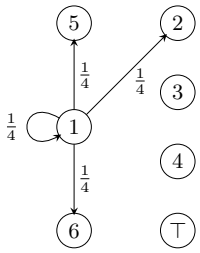

Fig. 5. Transition out of 1 in DTMC $\mathcal{D}$

Example 4. Consider an MDP $\mathcal{A}$ shown in Figure 2 where state 1 has two transitions one shown as solid edges and the other as dashed edges; transitions out of other states are not shown since they will not play a role. Suppose the atomic proposition a is $\top$ in $\{1, 2\}$ and $\bot$ in $\{3, 4\}$, and the proposition b is $\top$ in $\{1, 3\}$ and $\bot$ in $\{2, 4\}$. The formula $\varphi = \mathcal{P}_{\leq \frac{3}{4}}(Xa)$ evaluates to $\top$ in state 1.

The AMC $\mathcal{M}$ as defined in Definition 4, is shown in Figure 3. Now, because the distribution ν, given by $\nu(\{1\}) = \frac{1}{2}$, $\nu(\{2\}) = \frac{1}{2}$, $\nu(\{3\}) = 0$, and $\nu(\{4\}) = 0$ satisfies the bound constraints out of 1 but violates the property φ, φ evaluates to ? in state 1 of $\mathcal{M}$.

Now consider the tree-like partial order shown in Figure 4; arrows in the figure point from the smaller element to the larger one. If we construct the DTMC $\mathcal{D}$ over this partial order as in Definition 5, the transition out of state 1 will be as shown in Figure 5. Observe also that proposition a is $\top$ in $\{1, 2, 5\}$, $\bot$ in $\{3, 4, 6\}$ and ? in state $\top$; and proposition b is $\top$ in $\{1, 3\}$, $\bot$ in $\{2, 4\}$ and ? in $\{5, 6, \top\}$. Now φ evaluates to $\top$ in state 1, because the measure of paths out of 1 that satisfy $X\neg a$ is $\frac{1}{4}$. Thus, by Theorem 5, $\mathcal{M}$ is not simulated by $\mathcal{D}$. It is useful to observe that the upper bound managed to capture the constraint that the probability of transitioning to either 3 or 4 from 1 is at least $\frac{1}{4}$. Constraints of this kind that relate to the probability of transitioning to a set of states, cannot be captured by the interval constraints of an AMC, but can be captured by upper bounds on appropriate partial orders.

4 Abstracting CTMCs

We now outline how our upper bound construction gives us a way to abstract CTMC by other CTMCs. We begin with recalling the definitions of CTMCs, simulation and logical preservation, before presenting our abstraction scheme.

4.1 Preliminaries

The formulas of CSL are built up over a finite set of atomic propositions AP and are inductively defined as follows.

$$\varphi ::= \text{true} \mid a \mid \neg\varphi \mid \varphi \wedge \varphi \mid \mathcal{P}_{\bowtie p}(\varphi \, \mathcal{U}^t \, \varphi)$$

where $a \in \text{AP}$, $\bowtie \in \{<, \leq, >, \geq\}$, $p \in [0,1]$, and t a positive real number.

The 3-valued semantics of CSL is defined over *Continuous Time Markov Chains* (CTMC), where in each state every atomic proposition gets a truth value in $\mathbb{B}_3$. Formally, let Q be a finite set of states and let $\mathcal{Q} = (Q, \mathcal{P}(Q))$ be a measure space. A (uniform) CTMC [6] $\mathcal{M}$ is a tuple $(Q, \rightarrow, L, E)$, where $\rightarrow : Q \rightarrow \mathcal{M}_{=1}(\mathcal{Q})$, $L : (Q \times \text{AP}) \rightarrow \mathbb{B}_3$ is a labeling function that assigns a value in $\mathbb{B}_3$ to each atomic proposition in each state, and $E \in \mathbb{R}_{\geq 0}$ is the *exit rate* from any state. We will say $q \rightarrow \mu$ to mean $(q, \mu) \in \rightarrow$. Due to lack of space the formal semantics of the CTMC is skipped; the reader is referred to [18].

CSL's 3-valued semantics associates a truth value in $\mathbb{B}_3$ for each formula φ in a state q of the CTMC; we denote this by $[\![q, \varphi]\!]_{\mathcal{M}}$. The formal semantics is skipped and can be found in [13]. The model checking algorithm presented in [1] for the 2-valued semantics, can be adapted to the 3-valued case.

Simulation for uniform CTMCs, originally presented in [3], has been adapted to the 3-valued setting in [13] and is defined in exactly the same way as simulation in a DTMC; since the exit rate is uniform, it does not play a role. Once again, we say q_1 is simulated by q_2, denoted as $q_1 \preceq q_2$, iff there is a simulation $\sqsubseteq$ such that $q_1 \sqsubseteq q_2$. Once again, there is a close correspondence between simulation and the satisfaction of CSL formulas according to the 3-valued interpretation.

Theorem 6 (Katoen-Klink-Leucker-Wolf [13]). *Consider any states q, q' of CTMC $\mathcal{M}$ such that $q \preceq q'$. For any formula φ, if $[\![q', \varphi]\!]_{\mathcal{M}} \neq ?$ then $[\![q, \varphi]\!]_{\mathcal{M}} = [\![q', \varphi]\!]_{\mathcal{M}}$.*

4.2 Abstracting based on Upper Bounds

Abstraction can, once again, be accomplished by collapsing concrete states into a single abstract state on the basis of an equivalence relation on concrete states. The transition rates out of a single state can either be approximated by intervals giving upper and lower bounds, as suggested in [13], or by upper bound measures as we propose. Here we first present the proposal of Abstract CTMCs, where transition rates are approximated by intervals, before presenting our proposal. We conclude with a comparison of the two approaches.

[6] We only look at uniform CTMCs here; in general, any CTMC can be transformed in a uniform one that is weakly bisimilar to the original CTMC.

Definition 6. *Consider a CTMC* $\mathcal{M} = (Q_{\mathcal{M}}, \rightarrow_{\mathcal{M}}, L_{\mathcal{M}}, E_{\mathcal{M}})$ *with an equivalence relation* $\equiv$ *on* $Q_{\mathcal{M}}$. *An* Abstract CTMC *(ACTMC) [13] that abstracts* $\mathcal{M}$ *is a tuple* $\mathcal{A} = (Q_{\mathcal{A}}, \rightarrow_{\ell}, \rightarrow_{u}, L_{\mathcal{A}}, E_{\mathcal{A}})$, *where*

- $Q_{\mathcal{A}} = \{[q] \mid q \in Q_{\mathcal{M}}\}$ *is the set of equivalence classes of* $\equiv$,
- $E_{\mathcal{A}} = E_{\mathcal{M}}$,
- *If for all* $q_1, q_2 \in [q]$, $L_{\mathcal{M}}(q_1, a) = L_{\mathcal{M}}(q_2, a)$ *then* $L_{\mathcal{A}}([q], a) = L_{\mathcal{M}}(q, a)$. *Otherwise,* $L_{\mathcal{A}}([q], a) = ?$,
- $\rightarrow_{\ell} : Q_{\mathcal{A}} \rightarrow (Q_{\mathcal{A}} \rightarrow [0,1])$ *where*

$$[q] \rightarrow_{\ell} \mu \text{ iff } \forall [q_1] \in Q_{\mathcal{A}}\ \mu([q_1]) = \min_{q' \in [q]\ \wedge\ q' \rightarrow_{\mathcal{A}} \nu} \nu([q_1])$$

- *Similarly,* $\rightarrow_{u} : Q_{\mathcal{A}} \rightarrow (Q_{\mathcal{A}} \rightarrow [0,1])$ *where*

$$[q] \rightarrow_{u} \mu \text{ iff } \forall [q_1] \in Q_{\mathcal{A}}\ \mu([q_1]) = \max_{q' \in [q]\ \wedge\ q' \rightarrow_{\mathcal{A}} \nu} \nu([q_1])$$

Semantically, at a state $[q]$, *the ACTMC can make a transition according to any transition rates that satisfy the lower and upper bounds defined by* $\rightarrow_{\ell}$ *and* $\rightarrow_{u}$, *respectively.*

Katoen et al. demonstrate that the ACTMC $\mathcal{A}$ (defined above) does indeed simulate $\mathcal{M}$, and using Theorem 6 the model checking results of $\mathcal{A}$ can be reflected to $\mathcal{M}$. The measure of paths reaching a set of states within a time bound t can be approximated using ideas from [2], and this can be used to answer model checking question for the ACTMC (actually, the path measures can only be calculated upto an error).

Like in Section 3.2, we will now show how the upper bound construction can be used to construct (standard) CTMC models that abstract the concrete system. Before presenting this construction, it is useful to define how to lift a measure on a set with an equivalence relation $\equiv$, to a measure on the equivalence classes of $\equiv$.

Definition 7. *Given a measure* μ *on* $(Q, \mathcal{P}(Q))$ *and equivalence* $\equiv$ *on* Q, *the lifting of* μ *(denoted by* $[\mu]$*) to the set of equivalence classes of* Q *is defined as* $[\mu](\{[q]\}) = \mu(\{q' \mid q' \equiv q\})$.

Definition 8. *Let* $\mathcal{M} = (Q_{\mathcal{M}}, \rightarrow_{\mathcal{M}}, L_{\mathcal{M}}, E_{\mathcal{M}})$ *be a CTMC with an equivalence relation* $\equiv$ *on* $Q_{\mathcal{M}}$. *Let* $(Q, \sqsubseteq)$ *be a tree-like partial order with* $\top$, *such that* $\text{minimal}(Q) = \{[q] \mid q \in Q_{\mathcal{M}}\}$. *Let* $\mathcal{Q} = (Q, \mathcal{P}(Q), \sqsubseteq)$ *be the ordered measurable space over* Q. *Define the CTMC* $\mathcal{C} = (Q_{\mathcal{C}}, \rightarrow_{\mathcal{C}}, L_{\mathcal{C}}, E_{\mathcal{C}})$, *where*

- $Q_{\mathcal{C}} = Q$,
- $E_{\mathcal{C}} = E_{\mathcal{M}}$,
- *For* $q \in Q_{\mathcal{C}}$, *let* $\Gamma_q = \{[\mu] \mid \exists q' \in Q_{\mathcal{A}}.\ [q'] \sqsubseteq q \text{ and } q' \rightarrow_{\mathcal{A}} \mu\}$. *Now,* $q \rightarrow_{\mathcal{C}} \nabla(\Gamma_q)$, *and*
- *If for all* $q_1, q_2 \in Q_{\mathcal{M}}$ *such that* $[q_1] \sqsubseteq q$ *and* $[q_2] \sqsubseteq q$, $L_{\mathcal{M}}(q_1, a) = L_{\mathcal{M}}(q_2, a)$ *then* $L_{\mathcal{C}}(q, a) = L_{\mathcal{M}}(q_1, a)$. *Otherwise,* $L_{\mathcal{C}}(q, a) = ?$.

Once again, from the properties of least upper bounds, and definition of simulation, we can state and prove results analogous to Propositions 2 and 3. That is the CTMC $\mathcal{C}$ does indeed abstract $\mathcal{M}$ and it is the best possible on a given state space; the formal statements and proofs are skipped in the interests of space.

Comparison with Abstract CTMCs. All the points made when comparing the DTMC abstraction with the AMC abstraction scheme, also apply here. That is, the size of $\mathcal{C}$ is exponentially smaller than the size of the ACTMC $\mathcal{A}$. Moreover, we can choose the tree-like partial order used in the construction of $\mathcal{C}$ through a process of abstraction refinement. And finally, Example 4 can be modified to demonstrate that there are situations where the CTMC $\mathcal{C}$ gives a more precise result than the ACTMC $\mathcal{A}$. However, in the context of CTMCs there is one further advantage. ACTMCs can only be model checked approximately, while CTMCs can be model checked exactly. While it is not clear how significant this might be in practice, from a theoretical point of view, it is definitely appealing.

5 Conclusions

Our main technical contribution is the construction of least upper bounds for probability measures on measure spaces equipped with a partial order. We have developed an exact characterization of underlying orderings for which the induced ordering on probability measures admits the existence of least upper bounds, and provided a natural construction for defining them. We showed how these upper bound constructions can be used to abstract DTMCs, MDPs, and CTMCs by models that are purely probabilistic. In some situations, our abstract models yield more precise model checking results than previous proposals for abstraction. Finally, we believe that the absence of nondeterminism in the abstract models we construct might make their model-checking more feasible.

In terms of future work, it would be important to evaluate how these abstraction techniques perform in practice. In particular, the technique of identifying the right tree-like state space for the abstract models using abstraction-refinement needs to be examined further.

References

1. A. Aziz, K. Sanwal, V. Singhal, and R. Brayton. Model checking continuous-time Markov chains. *ACM TOCL*, 1:162–170, 2000.
2. C. Baier, B. Haverkrot, H. Hermanns, and J.-P. Katoen. Efficient computation of time-bounded reachability probabilities in uniform continuous-time Markov decision processes. In *Proc. of TACAS*, pages 61–76, 2004.
3. C. Baier, H. Hermanns, J.-P. Katoen, and V. Wolf. Comparative branching-time semantics for markov chains. *Inf. and Comp.*, 200:149–214, 2005.
4. R. Chadha, M. Viswanathan, and R. Viswanathan. Least upper bounds for probability measures and their applications to abstractions. Technical Report UIUCDCS-R-2008-2973, UIUC, 2008.

5. P. R. D'Argenio, B. Jeannet, H. E. Jensen, and K. G. Larsen. Reachability analysis of probabilistic systems by successive refinements. In *Proc. of PROBMIV*, pages 39–56, 2001.
6. P. R. D'Argenio, B. Jeannet, H. E. Jensen, and K. G. Larsen. Reduction and refinement strategies for probabilistic analysis. In *Proc. of PROBMIV*, pages 57–76, 2002.
7. J. Desharnais. *Labelled Markov Processes*. PhD thesis, McGill University, 1999.
8. H. Fecher, M. Leucker, and V. Wolf. Don't know in probabilistic systems. In *Proc. of SPIN*, pages 71–88, 2006.
9. M. Huth. An abstraction framework for mixed non-deterministic and probabilistic systems. In *Validation of Stochastic Systems: A Guide to Current Research*, pages 419–444. 2004.
10. M. Huth. On finite-state approximants for probabilistic computation tree logic. *TCS*, 346:113–134, 2005.
11. C. Jones. *Probabilistic Non-determinism*. PhD thesis, University of Edinburgh, 1990.
12. B. Jonsson and K. G. Larsen. Specification and refinement of probabilistic processes. In *Proc. of LICS*, pages 266–277, 1991.
13. J.-P. Katoen, D. Klink, M. Leucker, and V. Wolf. Three-valued abstraction for continuous-time markov chains. In *Proc. of CAV*, pages 311–324, 2007.
14. M. Kwiatkowska, G. Norman, and D. Parker. Game-based abstraction for markov decision processes. In *Proc. of QEST*, pages 157–166, 2006.
15. A. McIver and C. Morgan. *Abstraction, Refinement and Proof for Probabilistic Systems*. Springer, 2004.
16. D. Monniaux. Abstract interpretation of programs as markov decision processes. *Science of Computer Programming*, 58:179–205, 2005.
17. G. Norman. Analyzing randomized distributed algorithms. In *Validation of Stochastic Systems: A Guide to Current Research*, pages 384–418. 2004.
18. J. M. Rutten, M. Kwiatkowska, G. Norman, and D. Parker. *Mathematical Techniques for Analyzing Concurrent and Probabilistic Systems*. AMS, 2004.
19. N. Saheb-Djahromi. Probabilistic LCF. In *Proc. of MFCS*, pages 442–451, 1978.
20. R. Segala. Probability and nondeterminism in operational models of concurrency. In *Proc. of CONCUR*, pages 64–78, 2006.
21. K. Sen, M. Viswanathan, and G. Agha. Model checking markov chains in the presence of uncertainties. In *Proc. of TACAS*, pages 394–410, 2006.
22. B. Wachter, L. Zhang, and H. Hermanns. Probabilistic model checking modulo theories. In *Proc. of QEST*, 2007.
23. H. Younes, M. Kwiatkowska, G. Norman, and D. Parker. Numerical vs. statistical probabilistic model checking: An empirical study. In *Proc. of TACAS*, pages 46–60, 2004.
24. H. Younes and R. Simmons. Probabilistic verification of discrete event systems using acceptance sampling. In *Proc. of CAV*, pages 223–235, 2002.