A Branching-Time Theory for Probabilistic Model Checking*

S. Purushothaman Iyer[1], Rance Cleaveland[2], and Murali Narasimha[3]

[1] Department of Computer Science, North Carolina State University, Raleigh, NC 27695, USA
[2] Department of Computer Science, SUNY at Stony Brook, Stony Brook, NY 11794-4400, USA
[3] Wireless R&D Division, Ericsson, RTP, NC

NCSU CS Tech Report TR-2000-04
April 12, 2000

Abstract. This paper develops an approach to branching-time probabilistic model checking that is based on quantifying the likelihood with which a system satisfies a formula in the traditional mu-calculus. The semantics uniformly extends the standard interpretation of the mu-calculus and also subsumes work done in more restrictive probabilistic models. We also show how in our setting model checking may be reduced to equation-solving when the system in question is finite-state.

1 Introduction

The temporal-logic model-checking problem may be phrased as follows: given a system modeled as a transition system and a requirement formulated in temporal logic, does the system satisfy the formula? In traditional model-checking, system models can contain nondeterminism, and different schools of thought have arisen regarding the interpretation of temporal formulas over such systems. The *linear-time* view holds systems should be viewed as sets of sequences of states; in other words, the choices are "internal" and hence cannot be resolved, or affected by, and external observer of the system. The *branching-time* school treats systems semantically as trees whose nodes are states and whose edges represent transitions; the intuition is that choices are "external" and hence observable by the system's environment. Temporal logics reflecting both points of view have been developed, and various unifying logical theories, including CTL* and the mu-calculus, have been defined and studied, both from a theoretical and a practical standpoint.

For some classes of system properties, particularly those involving reliability and performance, one needs system models that include *probabilistic* information so that one can quantify the relative likelihoods of different kinds of operational

* Research supported by AFOSR grant F49620-95-1-0508, ARO grant P-38682-MA and NSF grants CCR-9505562, CCR-9996086 and INT-9996095.

behavior. An appealing paradigm for describing properties of such systems would then involve the use of a temporal logic to describe "desired behaviors" and a probabilistic semantics of the logic to enable reasoning about "how many" system executions satisfy the behavior. This paradigm has been advanced most notably for linear-time temporal logics [CY88,Var85], with work in branching-time probabilistic temporal logics tending to introduce constructs into the logic or its semantics to remove external choices, thereby "reducing" the problem to a linear-time one [PZ93,Han94,ASB$^+$95,BdA95b]. An interesting exception may be found in [HK97], which provides a probabilistic semantics of the mu-calculus, although one may criticize their approach for its ad hoc treatment of certain logical constructs.

In this paper we wish to advance a branching-time theory of temporal logics in which formulas in traditional logics are given a direct probabilistic interpretation. In pursuit of this goal, we consider a particular model of probabilistic labeled transition systems and a very expressive logic, the modal mu-calculus, and show how the logic may be interpreted in terms of measurable sets of "outcomes" in a traditional probabilistic manner. The main contributions of the paper are as follows.

- A semantics for a model of probabilistic labelled transition systems (NPLTS) freely incorporating nondeterministic and probabilistic choice.
- An interpretation of the modal mu-calculus against NPLTS that consistently extends interpretations of purely nondeterministic and purely probabilistic systems with respect to other temporal logics.
- A model-checking algorithm for calculating the probability with which an NPLTS system satisfies a mu-calculus formula.

The strength of our approach lies with the fact that branching-time assertions such as "both action a and action b are simultaneously enabled" are given an interpretation that is consistent with the one found in the traditional semantics. In contrast, the approaches based on Markov Decision Processes [Der70] as taken in [Han94,Seg95,BdA95a] would deem such a property to be impossible to satisfy.

2 Nondeterministic Probabilistic Labeled Transition Systems

In the remainder of the paper we fix two countable sets *Act* of *actions* and *Prop* of *atomic propositions*. Intuitively, *Act* contains the names of atomic actions that systems may engage in, while elements of *Prop* represent represent basic properties that may be used to characterize system states.

Mathematically, we model systems as labeled transition systems containing states that may be either *probabilistic* or *nondeterministic*. In the former transitions emanating from a state are selected on the basis of a coin flip, while in the latter case transitions from a state may be selected nondeterministically. The intention is that a system at a nondeterministic state may select its next transition nondeterministically, perhaps in response to a stimulus from the external world.

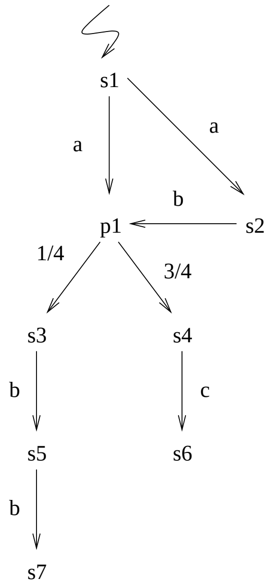

Fig. 1. A probabilistic nondeterministic labeled transition system (NPLTS).

Definition 1 (NPLTS). *A* nondeterministic probabilistic labeled transition system *(NPLTS) is a tuple $(S_N, S_P, \delta, \mathcal{P}, I)$ whose components are interpreted as follows.*

1. *$(n, n_1, \ldots \in) S_N$ is a countable set of nondeterministic states.*
2. *$(p, p_1, \ldots \in) S_P$ is a countable set of probabilistic states satisfying $S_N \cap S_P = \emptyset$. We use $(s, s_1, s', \ldots \in) S = S_N \cup S_P$ to represent the set of all states.*
3. *$\delta \subseteq S_N \times Act \times S$ is the nondeterministic transition relation, which is required to be* image-finite*: for every $n \in S_N$ and $a \in Act$, the set $\{s' \mid (n, a, s') \in \delta\}$ must be finite.*
4. *$\mathcal{P} : (S_P \times S_N) \to [0,1]$, the probabilistic state distribution, obeys $\sum_{n \in S_N} \mathcal{P}(p, n) = 1$ all $p \in S_P$.*
5. *$I : Prop \to 2^{S_N}$ is the* interpretation *function.*

Intuitively, a NPLTS describes the operational behavior of a system. The sets S_N and S_P contain the nondeterministic and probabilistic states of the system respectively. Relation δ records the execution steps enabled in nondeterministic states; such a step is labeled with an action from Act, with the intuition being that the system may engage in the transition when the action is enabled by the environment. Function $\mathcal{P}$ gives the probability distributions for the choices at each probabilistic choice state of the system. Note that every successor of a probabilistic state p is a nondeterministic state. We write $n \xrightarrow{a} s$ if $(n, a, s) \in \delta$ and $p \xrightarrow{r} n$ if $r > 0$ and $\mathcal{P}(p, n) = r$. In this notational scheme transitions can be labeled by elements of Act or $(0, 1]$; in the sequel we use $q, q_1, \ldots$ to range over the set $Act \cup (0, 1]$ of these possible transition labels. Figure 1 contains a graphical rendering of a NPTLS. $I(A)$ records the nondeterministic states for which A is true; note that in our framework, probabilistic states are treated as "ephemeral", or unobservable, and hence will not (directly) satisfy atomic propositions.

NPLTSs were first introduced by Hansson [Han94] as a general model of probabilistic computation. Other state-machine-oriented models of probabilistic computation have natural encodings as NPTLSs, and we briefly review these here. *Markov decision processes* [Der70] (MDPs) contain a mixture of probabilistic and nondeterministic states, although there is no notion of action. Consequently, MDPs may be represented as MLTSs whose action set $Act = \{a\}$ is a singleton set. *Markov chains* contain only probabilistic transitions; these may be represented as MLTSs with $Act = \{a\}$ and with δ satisfying the property that $\delta \subseteq S_N \times Act \times S_P$ is a function of type $S_N \times Act \rightarrow S_P$. *Reactive probabilistic labeled transition systems* [LS91] may be encoded as NPLTSs in which $\delta \subseteq S_N \times Act \times S_P$ is a partial function over $S_N \times Act$ that returns only probabilistic states. Standard (i.e. purely nondeterministic) labeled transition systems may be obtained by taking the set S_P of probabilistic states to be empty.

3 A Probabilistic Semantics of the Modal Mu-Calculus

This section presents a probabilistic semantics of the mu-calculus with respect to NPLTS. The crux of our probabilistic treatment of the mu-calculus is an interpretation of NPLTSs as "experiments", in the traditional probabilistic sense, whose "outcomes" are standard nondeterministic labeled transition systems. The standard semantics of the mu-calculus may be applied to these outcomes; we then wish to interpret the probability that the NPTLS satisfies a formula as the "measure", of the outcomes that satisfy the formula. This section formalizes these intuitions.

3.1 NPLTS as Experiments

We treat an NPLTS as probabilistic experiments by first unrolling the NPLTS into a (potentially infinite) tree; outcomes are then subtrees obtained by eliminating all but one probabilistic transition from each occurrence of probabilistic state. Figure 2 illustrates these ideas; in the specific example considered there all outcomes are finite trees because the NPLTS contains no loops, but this need not be the case in general. To formalize these notions we follow the construction laid out for reactive probabilistic labeled transition systems in [NCI99]. Specifically, we define the outcomes of a NPLTS as trees whose nodes are labeled by execution sequences, or "partial computations", of the NPLTS.

Definition 2. *Let $L = (S_N, S_P, \delta, \mathcal{P}, I)$ be a NPLTS. A sequence $s_0 \stackrel{q_1}{\rightarrow} s_1 \cdots \stackrel{q_k}{\rightarrow} s_k$, where $0 \leq k \leq \infty$, is a* partial computation *of L if for all $0 \leq i < k$, $s_i \stackrel{q_{i+1}}{\rightarrow} s_{i+1}$. We use $\mathcal{C}_L$ for the set of all partial computations of L and $\mathcal{C}_L(s)$ for all the partial computations whose initial state is s.*

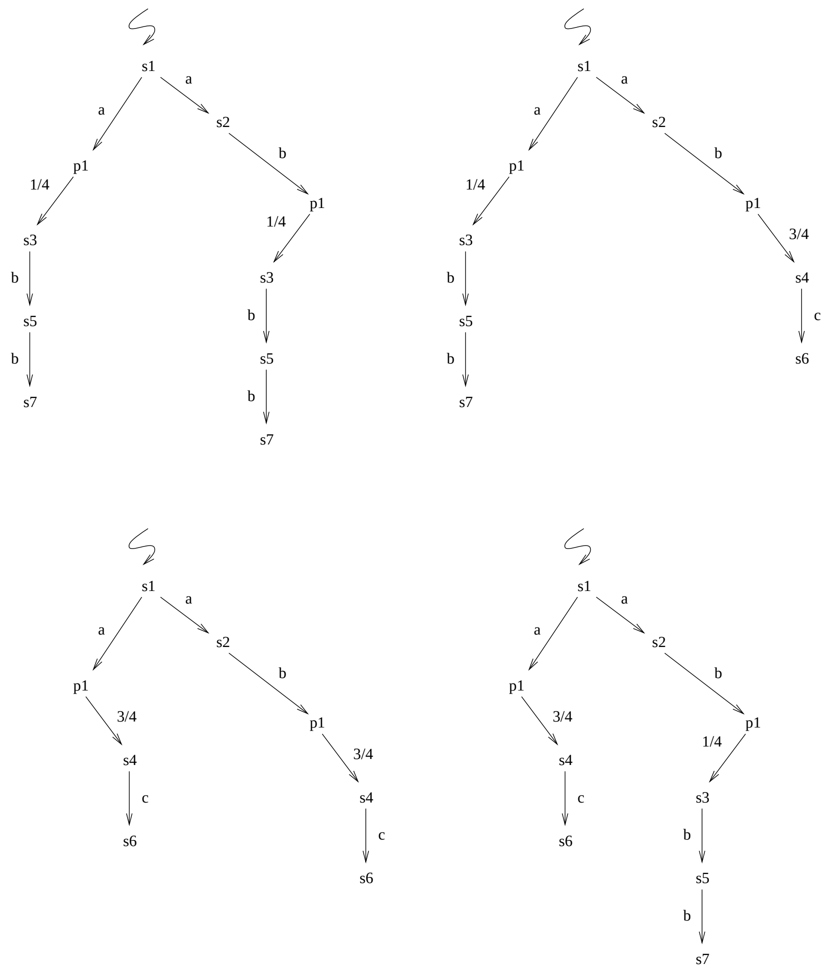

Fig. 2. The outcomes of the NPLTS in Figure 1.

The following operations and terminology on partial computations are used below.

Definition 3. *Let $L = (S_N, S_P, \delta, \mathcal{P}, I)$ be a NPLTS, let $\sigma = s_0 \xrightarrow{q_1} s_1 \cdots \xrightarrow{q_k} s_k$ and $\sigma' = s'_0 \xrightarrow{q'_1} s'_1 \cdots \xrightarrow{q_{k'}} s'_{k'}$ be elements of $\mathcal{C}_L$, and let $C \subseteq \mathcal{C}_L$.*

1. *If $k < \infty$ then σ is* finite, *and in this case we define $\mathsf{lst}(\sigma) = s_k$.*
2. *If $k < \infty$ and $s_k \xrightarrow{q} s'_0$ then $\sigma \xrightarrow{q} \sigma' \in \mathcal{C}_L$ is defined as $s_0 \xrightarrow{q_1} s_1 \cdots \xrightarrow{q_k} s_k \xrightarrow{q} s'_0 \xrightarrow{q'_1} s'_1 \cdots \xrightarrow{q_{k'}} s'_{k'}$.*
3. *C is* observable *if for every $\sigma, \sigma' \in C$, with $\sigma = s_0 \xrightarrow{q_1} s_1 \ldots s_k \xrightarrow{q_k} p \xrightarrow{r} n \ldots$ and $\sigma' = s_0 \xrightarrow{q_1} s_1 \ldots s_k \xrightarrow{q_k} p \xrightarrow{r'} n' \ldots$, then $n = n'$ and $r = r'$. (Recall that $p \in S_P$ and $n \in S_N$.)*

We also adopt the usual definitions of prefix *and* prefix-closed.

The notion of *observability* deserves further comment here. Roughly speaking, a set of partial computations is observable if any two distinct elements of the set differ only because they exercise different nondeterministic transitions.

The *observation trees* (o-trees for short) of a NPTLS are trees whose nodes correspond to partial computations and edges to transitions.

Definition 4. *Let $L = (S_N, S_P, \delta, \mathcal{P}, I)$ be a NPLTS. Then $\emptyset \neq T \subseteq \mathcal{C}_L$ is an* observation tree *if T is prefix-closed and observable, and if there exists an $s \in S$ such that $T \subseteq \mathcal{C}_L(s)$.*

If T is an o-tree then we use $\mathsf{root}(T)$ to refer to the s such that $T \subseteq C_L(s)$. (Note that in this case s is the only element in $S \cap T$.) We use $\mathcal{T}_L$ to refer to all o-trees of L and $\mathcal{T}_L(s) = \{T \in \mathcal{T}_L \mid \mathsf{root}(T) = s\}$. We say T' is a *prefix* of T if $T' \subseteq T$, and we call T is *finite* if $\mid T \mid < \infty$. T is *maximal* if there exists no o-tree T' with $T \subset T'$. We use $\mathcal{M}_L$ and $\mathcal{M}_L(s)$ to refer to the set of all maximal o-trees of L and the set of all maximal o-trees rooted at s respectively. It is precisely the set of maximal o-trees of L that we take to be the "outcomes" of the probabilistic experiment defined by L.

We close this treatment of outcomes and o-trees by defining the notion of subtree: intuitively, $T \xrightarrow{a} T'$ if T' is a subtree of T reachable via an edge labeled by a.

Definition 5. *Let $T, T' \in \mathcal{T}_L$. Then $T \xrightarrow{a} T'$ if at least one of the following holds.*

1. $\mathsf{root}(T) \in S$ *and* $\{\mathsf{root}(T) \xrightarrow{a} \sigma' \mid \sigma' \in T'\} \subseteq T$.
2. $\mathsf{root}(T) \in P$ *and for some* $r \in (0,1]$, $\{\mathsf{root}(T) \xrightarrow{r} s \xrightarrow{a} \sigma'' \mid \sigma'' \in T'\} \subseteq T$.

This definition reflects an essential intuition about NPTLS: probabilistic "transitions" are not execution steps in the same way that transitions labeled by actions are.

A measure space of outcomes. We now briefly outline how a measure space over $\mathcal{T}_L(s)$ may be defined for any state s in an NPLTS L. The construction is a straightforward extension of the standard sequence-space construction for Markov chains in which trees, rather than sequences, are the elements of the measurable sets. In what follows, fix $L = (S_N, S_P, \delta, \mathcal{P}, I)$ and $s \in S$.

Definition 6. *Let $T \in \mathcal{T}_L(s)$ be finite. Then the* basic cylindrical set *of T, $B_T \subseteq \mathcal{M}_L(s)$, is defined by $B_T = \{T' \in \mathcal{M}_L \mid T \subseteq T'\}$.*

We now define the measure of basic cylindrical sets the product of all probabilistic transitions in T.

Definition 7. *Let $T \in \mathcal{T}_L(s)$ be finite. Then the measure, $\mathsf{m}(B_T)$, of B_T is defined as $\mathsf{m}(B_T) = \prod_{\{r \in (0,1] \mid \sigma \xrightarrow{r} n \in T \,\wedge \mathsf{lst}(\sigma) \in S_P\}} r$.*

Intuitively, $\mathsf{m}(B_T)$ represents the proportion of all outcomes emanating from the root of B_T that have B_T as a prefix.

Consider the collection of basic cylindrical sets $\mathcal{B}^-(s)$ consisting of sets of the form B_T for finite T with $\mathsf{root}(T) = s$. We can then define a probability space $(\mathcal{M}_L(s), \mathcal{B}(s), \mathsf{m}_s)$ as follows.

Definition 8. *$\mathcal{B}(s)$ is the smallest field of sets containing $\mathcal{B}^-(s)$ and closed with respect to denumerable unions and complementation. $\mathsf{m}_s : \mathcal{B}(s) \to [0,1]$ is then defined inductively as follows.*

$$\begin{aligned}
\mathsf{m}_s(B_T) &= \mathsf{m}(B_T) \\
\mathsf{m}_s(\bigcup_{i \in I} B_i) &= \sum_{i \in I} \mathsf{m}_u(B_i) \text{ for pairwise disjoint } B_i \\
\mathsf{m}_s(B^c) &= 1 - \mathsf{m}_s(B)
\end{aligned}$$

It is easy to show that m_s is a probability measure over $\mathcal{B}(s)$. Consequently, $(\mathcal{M}_L(s), \mathcal{B}(s), \mathsf{m}_s)$ is indeed a probability space. We refer to a set $M \subseteq \mathcal{M}_L(s)$ as *measurable* if $M \in \mathcal{B}(s)$.

3.2 Measurability and the Modal Mu-Calculus

We now introduce the modal mu-calculus and show how the previously defined measure space may be used to give a probabilistic interpretation to the alternation-free formulas of the logic. Fix a set $(X, Y\ in) Var$ of propositional variables. The syntax of the mu-calculus is given by the following grammar.

$$\psi ::= X \mid A \mid \neg A \mid \psi_1 \vee \psi_2 \mid \psi_1 \wedge \psi_2 \mid \langle a \rangle \psi \mid [a]\psi \mid \mu X.\psi \mid \nu X.\psi$$

We adopt the usual definitions of bound and free variables, open and closed terms, and subformula. The well-formed formulas of the mu-calculus are the closed formulas generated by the above rules; we use $(\psi, \psi', \psi_1, \ldots \in)\Psi$ to represent the set of such formulas. A formula $\psi \in \Psi$ is *alternation-free* if every free variable in every subformula of ψ of form $\mu X.\psi'$ is bound by μ in ψ, and similarly for ν.

Traditionally, the mu-calculus is interpreted with respect to states in a given labeled transition system: the formal semantics assigns to each formula a set of states for which the formula "holds" [Koz83]. The $\langle a \rangle$ and $[a]$ operators are action-labeled modalities; a state satisfies $\langle a \rangle \psi$ if it has an a-transition to a state satisfying ψ, while a state satisfies $[a]\psi$ if all its a-transitions have this property. The μ and ν constructs are *fixpoint* operators, with $\mu X.\psi$ representing the "smallest" set of states "satisfying" $X = \psi$ and $\nu X.\psi$ the largest such set.

In order to give a probabilistic semantics to the alternation-free mu-calculus we first show how, given a NPLTS L, any formula may be interpreted as sets of outcomes. Then we establish that these sets are measurable in the alternation-free case. For a formula ψ and environment e mapping variables to sets of outcomes, $\Theta_L(\psi)e$ returns the set of outcomes in L satisfying ψ; e is needed to handle free variables.

Definition 9. *Let $L = (S_N, S_P, \delta, \mathcal{P}, I)$ be a NPLTS. The function $\Theta_L : \Psi \to (Var \to 2^{\mathcal{M}_L}) \to 2^{\mathcal{M}_L}$ is defined as follows.*

- $\Theta_L(A)e = \bigcup_{n \in I(A)} \mathcal{M}_L(n) \;\cup\; \bigcup_{p \in S_P} \{T \in \mathcal{M}_L(p) \mid \exists n \in I(A), r \in (0,1].p \xrightarrow{r} n \in T\}$
- $\Theta_L(\neg A)e = \mathcal{M}_L - \Theta_L(A)e$
- $\Theta_L(X)e = e(X)$
- $\Theta_L(\psi_1 \vee \psi_2)e = \Theta_L(\psi_1) \cup \Theta_L(\psi_2)e$
- $\Theta_L(\psi_1 \wedge \psi_2)e = \Theta_L(\psi_1)e \cap \Theta_L(\psi_2)e$
- $\Theta_L(\langle a \rangle \psi)e = \{T \mid \exists T' : T \xrightarrow{a} T' \wedge T' \in \Theta_L(\psi)e\}$
- $\Theta_L([a]\psi)e = \{T \mid \forall T',\ (T \xrightarrow{a} T') \Rightarrow T' \in \Theta_L(\psi)e\}$
- $\Theta_L(\mu X.\psi)e = \bigcap\{M \mid \Theta_L(\psi)e[X \mapsto M] \subseteq M\}$
- $\Theta_L(\nu X.\psi)e = \bigcup\{M \mid M \subseteq \Theta_L(\psi)e[X \mapsto M]\}$

The semantics of atomic propositions A deserves some comment. A tree can "satisfy" A in one of two ways. Firstly, if the tree's root is a nondeterministic state, and that state satisfies A, then the tree does. Secondly, if the tree's root is a probabilistic state, and the (by definition sole) probabilistic transition leads to a nondeterministic state satisfying A, then the tree does as well. This reflects our desire to tree probabilistic states as "ephemeral". The other operators have standard interpretations taken from the mu-calculus. The meaning of the fixpoint operators μ and ν relies on the characterization of least and greatest fixpoints of monotonic functions over complete lattices (here, the lattice of sets of trees) due to Knaster and Tarski [Tar55]. One may easily show that a tree T is in $\Theta_L(\psi)e$ if and only if the start state of T is among the "states" of T satisfying ψ with respect to the usual mu-calculus semantics [Koz83]. Finally, if s is a state in L then we define $\Theta_{L,s}(\psi)e = \Theta_L(\psi)e \cap \mathcal{M}_L(s)$ to be the set of s-rooted trees satisfying ψ in e.

Measurability of $\Theta_{L,s}(\psi)e$. We now establish that for closed alternation-free formulas ψ and state s, the set $\Theta_{L,s}(\psi)e$ of outcomes rooted at s and satisfying ψ is measurable, and hence that $\mathsf{m}_s(\Theta_{L,s}(\psi)e)$ is well-defined. Intuitively, this quantity reflects the "probability" that system L, when in state s, satisfies ψ.

Theorem 1. *Let $L = (S_N, S_P, \delta, \mathcal{P}, I)$ be a NPLTS with $s \in S_N \cup S_P$, and let ψ be a closed alternation-free mu-calculus formula. Then $\mathsf{m}_s(\psi) = \Theta_{L,s}(\psi)e$ is measurable.*

Proof sketch The proof depends on the following observation.

> Let ψ be a μ- and ν-free formula with free variable X, and let e be an environment. Then the function $f : 2^{\mathcal{M}_L} \to 2^{\mathcal{M}_L}$ given by $f(M) = \Theta(\psi)e[X \mapsto M]$ is continuous.

This fact is a result of the continuity of the non-fixpoint operators in the logic; note that the continuity of $[a]$ depends on the image-finiteness of L.

This continuity result implies that the semantics of μ and ν may be given iteratively; in particular, if ψ is μ- and ν-free with free variable X then we have that $\Theta(\mu X.\psi)e = \bigcup_{i=0}^{\infty} M_i$ and $\Theta(\nu X.\psi)e = \bigcup_{i=0}^{\infty} \widehat{M}_i$, where $M_0 = \emptyset$ and

$M_{i+1} = \Theta(\psi)e[X \mapsto M_i]$, and $\widehat{M_0} = \mathcal{M}_L$ and $\widehat{M_{i+1}} = \Theta(\psi)e[X \mapsto \widehat{M_i}]$. It is easy to show that each M_i is measurable, and as countable unions of measurable sets are also measurable, the result follows (and dually for $\widehat{M_i}$). ■

In closing we remark that the above argument breaks down when non-alternation-free formulas are allowed, since in general, the μ and ν operators are not continuous. We nevertheless conjecture that that the full mu-calculus is indeed measurable.

4 Model Checking

This section develops a procedure for answering the following question: given a NPLTS with state s, and given an alternation-free mu-calculus formula ψ, what is $\mathsf{m}_s(\psi)$? The approach we take relies on (a) building a graph from s and ψ, and (b) deriving and solving equations from the graph. We require that ψ contain no unguarded variables: every free occurrence of a variable within ψ must fall within the scope of a modality. This restriction is not serious, as every mu-calculus formula can be transformed into one meeting this restriction.

The graph that we build will employ two kinds of nodes, *state assertions* and *transition assertions*. Both are interpreted semantically as sets of trees, with the former constraining trees emanating from a state and the latter imposing restrictions on the initial edges present in a tree.

Definition 10 (State assertion). *A* state assertion *is a tuple (s, F) where $s \in S$ and $F \subseteq \Psi$ is a set of closed formulae. The semantics of a state assertion $\gamma = (s, F)$ is given by: $[\![(s, F)]\!] = \mathcal{M}_L(s) \cap \Theta_L(\wedge F)$* ■

Definition 11 (Transition assertion). *A transition assertion $\rho \subseteq S_N \times \Psi \times S$ satisfies: (a) all triples (n, ψ, s) in γ have the same first component n, (b) the formula ψ in every triple is of the form $\langle a\rangle\psi$ or $[a]\psi$, and (c) for every $(n, \langle a\rangle\psi, s)$ or $(n, [a]\psi, s)$ in ρ there is a transition $n \xrightarrow{a} s$ in the NPLTS. The semantics of a triple is defined as: $[\![(n, \langle a\rangle\psi, s)]\!] = [\![(n, [a]\psi, s)]\!] = \mathsf{f}_{\langle a\rangle}(\Theta_{L,n}(\psi)) \cap \mathcal{M}_L(n)$, where $\mathsf{f}_{\langle a\rangle}(M) = \{T \in \mathcal{M}_L \mid \exists T' \in M.T \xrightarrow{a} T'\}$. The semantics of a transition assertion ρ is given as: $[\![\rho]\!] = \bigcap_{(n,\psi,s)\in\gamma}[\![(n, \psi, s)]\!]$* ■

We now introduce the following operations on transition assertions.

Definition 12. *Let ρ be a transition assertion. Then the* transitions *of ρ are defined as:*

$$trans(\rho) = \{n \xrightarrow{a} s \mid (n, \langle a\rangle\psi, s) \in \rho \vee (n, [a]\psi, s) \in \rho\}.$$

The residue *of ρ with respect to $n \xrightarrow{a} s$ be defined as:*

$$residue(\rho, n \xrightarrow{a} s) = \{\psi \mid (n, \langle a\rangle\psi, s) \in \rho \vee (n, [a]\psi, s) \in \rho\}$$

■

Given state s and closed alternation-free formula ψ, the graph we build will start with a node containing the state assertion $(s, \{\psi\})$ and will use the structure of both to add new assertions; this is to be expected given that wish to calculate $\mathsf{m}_s(\psi)$. Transition assertions arise as a result of processing *modal (state) assertions* – state assertions of the form (s, F) where all formula in F are of the form $\langle a\rangle\psi$ or $[a]\psi$. Since in general $[\![(n, F)]\!] = \Theta_{L,n}(\bigwedge F)$ we are interested in n-rooted trees that satisfy each of the formulas in F. If a tree T satisfies $[a]\psi'$ then each a-subtree T should satisfy ψ'. Consequently, there is only one way in which T could satisfy $[a]\psi'$. However, a formula of the form $\langle a\rangle\psi'$ could be satisfied by a tree T in more than one way, depending on which of the a-subtrees of T satisfies ψ'. In effect, we have a conjunction (due to (n, F)) of disjunctions (due to the $\langle\rangle$-formulae) giving rise to several transition assertions. We formalize this discussion as follows by first introducing when a transition assertion is *reasonable* with respect to a modal state assertion:

Definition 13. *Given a modal state assertion $\gamma = (n, F)$ define transition assertion ρ to be* reasonable *with respect to γ (written as* $\mathsf{Reasonable}(\gamma, \rho)$*) provided:*

- *All triples in ρ have form (n, ϕ, s) with $\phi \in F$, and $trans(\{(s, \psi, u)\}) \subseteq \delta$.*
- *For every formula $[a]\psi$ in F we have $\{(n, [a]\psi, s) | n \xrightarrow{a} s\} \subseteq \rho$, and for every formula $\langle a\rangle\psi$ in F we have $\{(n, \langle a\rangle\psi, s) | n \xrightarrow{a} s\} \cap \rho \neq \emptyset$.*

Let $\mathcal{R}(\gamma) = \{\rho | \mathsf{Reasonable}(\gamma, \rho)\}$. ■

Note that the definition of *reasonable*ness is semantics-preserving: if $\mathsf{Reasonable}(\gamma, \rho)$ then $[\![\rho]\!] \subseteq [\![\gamma]\!]$. Since a tree T could be justified to be in $[\![\gamma]\!]$ based on several $\rho \in \mathcal{R}(\gamma)$ we need the notion of $rank(\gamma, \rho)$ to describe whether $\mathsf{m}_s([\![\rho]\!])$ would appear positively or negatively in the probability equation for $\mathsf{m}_s([\![\gamma]\!])$.[1]

Definition 14. *Let $\gamma = (n, F)$ be a modal state assertion and let* $\mathsf{Reasonable}(\gamma, \rho)$*. Then $rank(\gamma, \rho) = \prod_{\langle a\rangle\psi \in F} odd(\{(n, \langle a\rangle\psi, s) \in \rho\})$ where $odd(A) = \begin{cases} +1 & \text{if } |A| \text{ is odd} \\ -1 & \text{otherwise} \end{cases}$*
■

We are now ready to present our graph construction.

Definition 15. *The graph* $\mathrm{Pr}(s, \psi) = (\Gamma, \Delta, \ell)$*, where*

- *$\Gamma \subseteq (S \times 2^{\Psi}) \cup 2^{S_N \times \Psi \times S}$ are the nodes containing either a state assertion or a transition assertion.*
- *$\Delta \subseteq \Gamma \times (Act \cup \{\epsilon\} \cup \{r \mid 0 < r \leq 1\}) \times \Gamma$ is the set of edges.*
- *$\ell : \Delta \to \{+1, -1\}$ is a labelling function.*

The sets of nodes, edges and their labellings are defined, inductively, as follows:

1. *$(s, \{\psi\}) \in \Gamma$.*

[1] Note that $prob(A \cap B)$ appears negatively in $prob(A \cup B) = prob(A) + prob(B) - prob(A \cap B)$.

2. *If* $\gamma = (p, F) \in \Gamma$ *some* $p \in S_P$ *then for each probabilistic edge* $p \xrightarrow{r} n$ *in* L *add a node* $\gamma' = (n, F)$ *to* Γ, *and an edge* (γ, r, γ') *to* Δ *and set* $\ell(\gamma, r, \gamma') = +1$.
3. *If* $\gamma = (n, F)$ *is a modal assertion then for each* $\rho \in \mathcal{R}(\gamma)$ *add* ρ *to* Γ *and* (γ, ϵ, ρ) *to* Δ *and set* $\ell(\gamma, \epsilon, \rho)$ *to* $rank(\gamma, \rho)$.
4. *If* $\gamma = (n, F)$ *is a non-modal assertion such that* $n \in S_N$ *then the first applicable rule among the following dictates what the successor nodes to* γ, *the edges from* γ *and the labels on the edges are.*
   (a) *If (1)* $\psi = A \in F$ *and* $n \in I(A)$, *or (2)*$\psi = \neg A \in F$ *and* $n \notin I(A)$, *or (3)*$\psi = [a]\psi' \in F$ *and* $n \not\xrightarrow{a}$ *then* $\gamma' = (n, F - \{\psi\})$, $(\gamma, \epsilon, \gamma') \in \Delta$ *and* $\ell(\gamma, \epsilon, \gamma') = +1$.
   (b) *If (1)* $\psi = \neg A \in F$ *and* $n \in I(A)$, *or (2)*$\psi = A \in F$ *and* $n \notin I(A)$, *or (3)*$\psi = \langle a\rangle\psi' \in F$ *and* $n \not\xrightarrow{a}$ *then* $\gamma' = (n, \{false\})$, $(\gamma, \epsilon, \gamma') \in \Delta$ *and* $\ell(\gamma, \epsilon, \gamma') = +1$.
   (c) *If* $\psi = \psi_1 \wedge \psi_2 \in F$ *then* $\gamma' = (n, (F - \{\psi\}) \cup \{\psi_1, \psi_2\})$, $(\gamma, \epsilon, \gamma') \in \Delta$, *and* $\ell(\gamma, \epsilon, \gamma') = +1$.
   (d) *If* $\psi = \psi_1 \vee \psi_2 \in F$ *then there are three successors to* γ. *They are (a)* $\gamma_1 = (n, (F - \{\psi\}) \cup \{\psi_1\})$, *(b)* $\gamma_2 = (n, (F - \{\psi\}) \cup \{\psi_2\})$ *and (c)* $\gamma_3 = (n, (F - \{\psi\}) \cup \{\psi_1 \wedge \psi_2\})$. *The three edges are* $(\gamma, \epsilon, \gamma_i)$ *for* $1 \leq i \leq 3$. *The edge labels are* $\ell(\gamma, \epsilon, \gamma_1) = \ell(\gamma, \epsilon, \gamma_2) = +1$ *and* $\ell(\gamma, \epsilon, \gamma_3) = -1$.
   (e) *If* $\psi = \nu X.\psi' \in F$ *or* $\psi = \mu X.\psi' \in F$ *then* $\gamma' = (s, (F - \{\psi\}) \cup \{\psi'[X \mapsto \psi]\})$, $(\gamma, \epsilon, \gamma') \in \Delta$ *and* $\ell(\gamma, \epsilon, \gamma') = 1$.
5. *If* γ *is a transition assertion with a common first component* $n \in S_N$ *then for every* $n \xrightarrow{a} s \in trans(\gamma)$ *add* $\gamma' = (n, residue(\gamma, n \xrightarrow{a} s))$ *to* Γ, (γ, a, γ') *to* Δ *and* $\ell(\gamma, a, \gamma') = +1$. ■

For any s and closed ψ $\Pr(s, \{\psi\})$ is finite due to the finiteness of the NPLTS and the fact that the formulae that arise in the construction are all in the Fisher-Ladner closure of ψ (which is a finite set). The relation among the sets of trees denoted by the nodes and the relation between their measures is given by the following:

Lemma 1. *Let* γ *be a node. Then the following hold:*

- *If* $\gamma = (p, F)$, *where* $p \in S_P$, *then* $[\![\gamma]\!] = \bigcup_{p \xrightarrow{r} n} \{p \xrightarrow{r} T \mid T \in [\![(n, F)]\!]\}$ *and* $\mathsf{m}_p([\![(p, F)]\!]) = \sum_{p \xrightarrow{a} n} r \times \mathsf{m}_n([\![(n, F)]\!])$
- *If* $\gamma = (n, F)$ *is a state assertion with successors* $\gamma_1, \ldots \gamma_k$ *then* $[\![\gamma]\!] = \bigcup_{i=1}^{k} [\![\gamma_i]\!]$ *and* $\mathsf{m}_n([\![\gamma]\!]) = \sum_{1 \leq i \leq k} \ell(\gamma, \epsilon, \gamma_i) \times \mathsf{m}_n([\![\gamma_i]\!])$
- *If* γ *is a transition assertion with a common first component* $n \in S_N$ *then*

$$[\![\gamma]\!] = \bigcap_{n \xrightarrow{a} s \in trans(\gamma)} \mathsf{f}_{\langle a,n\rangle}([\![(s, residue(\gamma, n \xrightarrow{a} s))]\!]) \tag{1}$$

$$\mathsf{m}_n([\![\gamma]\!]) = \prod_{n \xrightarrow{a} s \in trans(\gamma)} \mathsf{m}_s([\![(s, residue(\gamma, n \xrightarrow{a} s))]\!]) \tag{2}$$

where $\mathsf{f}_{\langle a,n\rangle}(M) = \mathsf{f}_{\langle a\rangle}(M) \cap \mathcal{M}_L(n)$.

Proof sketch: The first two cases are routine. The only hard case is when γ is a transition assertion, in which case the following argument based on distributivity of $\mathsf{f}_{\langle a,s\rangle}$ over intersection proves the relation between the sets of trees:

$$\begin{aligned}[\![\gamma]\!] &= \bigcap_{(n,\psi,s)\in\gamma}[\![(n,\psi,s)]\!]\\ &= \bigcap_{(n,\langle a\rangle\psi',s)\in\gamma}\mathsf{f}_{\langle a,n\rangle}(\Theta_{L,s}(\psi'))\cap\bigcap_{(n,[a]\psi',s)\in\gamma}\mathsf{f}_{\langle a,n\rangle}(\Theta_{L,s}(\psi'))\\ &= \bigcap_{n\xrightarrow{a}s\in trans(\gamma)}\mathsf{f}_{\langle a,n\rangle}(\Theta_{L,s}(\bigwedge_{\psi'\in residue(\gamma,n\xrightarrow{a}s)}\psi'))\end{aligned}$$

As regards the measure equations the first two cases are routine. The third depends upon the observations that (a) for two different transitions $n\stackrel{a_1}{\to}s_1$ and $n\stackrel{a_2}{\to}s_2$ the sets $\mathsf{f}_{\langle a_1,n\rangle}([\![(s_1,residue(\gamma,n\stackrel{a_1}{\to}s_1))]\!])$ and $\mathsf{f}_{\langle a_2,n\rangle}([\![(s_2,residue(\gamma,n\stackrel{a_2}{\to}s_2))]\!])$ are independent (as for every $T_2\in\mathcal{M}_{s_2}$ and every $T_1\in[\![(s_1,F_1)]\!]$ there is a tree $T\in\mathsf{f}_{\langle a_1,n\rangle}([\![(s_1,F_1)]\!])$ such that T_1 is a a_1-subtree of T and T_2 is an a_2-subtree of T) and (b) that for every transition $n\stackrel{a}{\to}s$ we have $\mathsf{m}_n(\mathsf{f}_{\langle a,n\rangle}(\Theta_{L,s}(\wedge residue(\gamma,n\stackrel{a}{\to}s)))) = \mathsf{m}_s(\Theta_{L,s}(\wedge residue(\gamma,n\stackrel{a}{\to}s)))$. ■

Calculating measures We now explain how to generate a system of equations from the graph described above. The system will contain one variable, X_γ, for each node γ in the graph and one equation containing this variable as its left-hand side. The right-hand side of the equation, E_γ, for X_γ is generated as follows:

$$E_\gamma=\begin{cases}0 & \text{if } \gamma=(s,\{false\}) \text{ for some } s\\ 1 & \text{if } \gamma=(s,\emptyset) \text{ for some } s\\ \sum_{(\gamma,r,\gamma')\in\Delta} r\cdot X_{\gamma'} & \text{if } \gamma=(p,F)\wedge p\in S_P\\ \prod_{(\gamma,a,\gamma')\in\Delta} X_{\gamma'} & \text{if } \gamma \text{ is a set of triples}\\ \sum_{(\gamma,\epsilon,\gamma')}\ell(\gamma,\epsilon,\gamma')\cdot X_{\gamma'} & \text{if } \gamma=(s,F)\wedge s\in S_N\end{cases}$$

Lemma 1 indicates that $\mathsf{m}_s([\![\gamma]\!])$, where $[\![\gamma]\!]\subseteq\mathcal{M}_L(s)$, is a solution to the equations generated. However, in general there will be many such solutions, and the question then arises as to how we determine which solution indeed corresponds to the measures we want. We suggest the following procedure:

1. Compute the strongly connected components of the graph from which the equations are constructed and topologically sort them.
2. Propagate solutions as far as possible: If a solution has been computed for a variable, replace all occurrences of the variable in the right-hand sides by the variable.
3. Beginning at the end of the strongly connected component list, process each component C as follows.
   (a) If C contains a μ-node (i.e. a state assertion containing at least one μ-formula), assign each variable corresponding to a node in C the value 0; otherwise, assign 1 to each variable.
   (b) Repeatedly calculate new values for the variables of C by evaluating each right-hand side using the old values. Stop when values do not change (or fall within a tolerance ϵ that is a parameter to the algorithm).
   (c) Propagate these values.

In general, this algorithm requires the specification of an "error tolerance" ϵ because the quantities being manipulated are real numbers. So the algorithm is approximation-based. However, all the functions being used are continuous, and hence the iteration process described above converges. We now have the following.

Lemma 2. *Let $s \in S$ and ψ be a closed formula. Then the quantity calculated for $X_{(s,\{\psi\})}$ converges to $\mathsf{m}_s(\psi)$.*

Comparison of systems of equations: Assume that we are given an NPLTS which has no nondeterminism, i.e., in all nonprobabilistic nodes there is at most one edge for each $a \in Act$. In such a cases, the set of equations we build will have the property that for every modal state assertion there is exactly one successor, which is a transition assertion. This is exactly the structure of equations we get for reactive systems [NCI99] where there is no non-deterministic choice. If in addition the set of actions is a singleton set (a model that corresponds to Markov processes), then an inspection of our equations indicates they will be linear – which is what is used in model-checking Markov chains [CY88] against LTL. Finally, if the NPLTS contains no probabilistic states then all solutions will either be 0 or 1 and will coincide with the standard nondeterministic semantics.

5 Conclusions

This paper has developed a probabilistic interpretation of the modal mu-calculus with respect to models containing both probabilistic and nondeterministic choice. The semantics relies on a formalization of a probability space of outcomes that these models naturally give rise to. Alternation-free mu-calculus formulas turn out to define measurable sets of such outcomes, and these measures may be characterized as solutions to sets of equations. As future work we wish to show that the full mu-calculus is measurable and to investigate fast techniques for solving the equations generated by our model-checking procedure. We also wish to look into verifying communications protocols with probabilistic behavior and to investigate other quantitative information on transition systems that may be computable using our equation-solving techniques.

References

[ASB+95] A. Aziz, V. Singhal, F. Balarin, R.K. Brayton, and A.L. Sangiovanni-Vincentelli. It usually works: The temporal logic of stochastic systems. In *Proc of Computer Aided Verification '95*. Springer-Verlag, July 1995.

[BdA95a] A. Bianco and L. de Alfaro. Model checking of probabilistic and mondeterministic systems. In *Foundations and Software Technology and Theoretical Computer Science*, volume 1026 of *LNCS*, pages 499–513. Springer-Verlag, 1995.

[BdA95b] A. Bianco and L. de Alfaro. Model-checking of probabilistic and non-deterministic systems. In *Proc. Foundations of software technology and theoretical computer science*, Lecture notes in Computer science, vol 1026, pages 499–513. Springer-Verlag, December 1995.

[CY88] C. Courcoubetis and M. Yannakakis. Verifying temporal properties of finite-state probabilistic programs. In *Proc. 1988 IEEE Symp. on the Foundations of Comp. Sci.*, 1988.

[Der70] C. Derman. *Finite State Markov Decision Processes*. Academic Press, 1970.

[Han94] H. Hansson. *Time and Probability in Formal Design of Distributed Systems.* Elsevier, 1994.

[HK97] Michael Huth and Marta Kwiatkowska. Quantitative analysis and model checking. In *Proceedings, Twelth Annual IEEE Symposium on Logic in Computer Science*, pages 111–122, Warsaw, Poland, 29 June–2 July 1997. IEEE Computer Society Press.

[Koz83] D. Kozen. Results on the propositional μ-calculus. *Theoretical Computer Science*, 27(1):333–354, 1983.

[LS91] K. G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94, 1991.

[NCI99] Murali Narasimha, Rance Cleaveland, and Purush Iyer. Probabilistic temporal logics via the modal mu-calculus. In Wolfgang Thomas, editor, *Proceedings of Foundations of Software Science and Computation Structures*, pages 288–305. ETAPS, Springer, March 1999.

[PZ93] Amir Pnueli and Lenore D. Zuck. Probabilistic verification. *Information and Computation*, 103(1):1–29, March 1993.

[Seg95] R. Segala. A compositional trace-based semantics for probabilistic automata. In *CONCUR95*, pages 324–338, 1995.

[Tar55] A. Tarski. A fixed point theorem and its applications. *Pacific J. Math.*, pages 285–309, 1955.

[Var85] M. Y. Vardi. Automatic verification of probabilistic concurrent finite-state programs. In *IEEE Symposium on Foundations of Computer Science*, pages 327–338, 1985.