Branching-Time Probabilistic Model Checking*

R. CLEAVELAND
SUNY at Stony Brook, USA

S. PURUSHOTHAMAN IYER
North Carolina State University, USA

Abstract

This paper presents a probabilistic semantics for the alternation-free modal mu-calculus. This semantics precisely defines the probability with which a system that includes a mixture of probabilistic and nondeterministic choices satisfies a formula in the logic. An algorithm for computing these probabilities is also given.

Keywords

model checking, branching time, temporal logic, mu-calculus, probabilistic state machines

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

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? The system models used in traditional model-checking typically contain nondeterminism, and different approaches have been proposed for defining temporal formulas over such systems. In *linear-time* logics systems are viewed as sets of sequences of states: nondeterministic choices are resolved "internally". On the other hand, *branching-time* logics treat systems as trees recording choice points explicitly; thus choices are "external" and observable by the system's environment. Temporal logics reflecting both points of view have been developed, and various unifying logical theories, including CTL* [5] and the mu-calculus [8], have been devised.

Many aspects of system behavior, particularly those involving reliability and performance, require system system models that include *probabilistic* information. An appealing paradigm for describing properties of such systems involves the use of temporal formulas to describe "desired behaviors" and a probabilistic semantics of the logic to enable reasoning about "how many" system executions are "desirable". This paradigm has been advanced most notably for linear-time temporal logics [3, 14]. Work in branching-time probabilistic temporal logics has tended toward the introduction constructs into the logic or its semantics to remove nondeterminism choices, thereby "reducing" the problem to a linear-time one [1, 2, 6, 11]. An interesting exception may be found in [7], which gives a probabilistic interpretation of the mu-calculus in terms of ranges of probabilities.

In this paper we wish to advance a branching-time theory of temporal logics in which formulas are given a direct and exact probabilistic interpretation. In pursuit of this goal, we consider a general model of probabilistic 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 measure-theoretic treatment of probabilistic labeled transition systems (NPLTS) freely incorporating nondeterministic and probabilistic choice.
- An interpretation of the modal mu-calculus against NPLTS that consistently extends those for purely nondeterministic and purely probabilistic systems.
- 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.

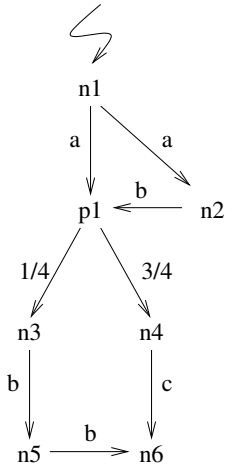

Figure 1: A nondeterministic probabilistic labeled transition system (NPLTS).

2 Labeled Transition Systems

In the remainder of the paper we fix two countable sets: *Act* contains the *atomic actions* systems may perform, while and *Prop* consists of the *atomic propositions* that may hold in system states.

We model systems as labeled transition systems containing both *probabilistic* and *nondeterministic* states. Transitions emanating from the former are selected on the basis of a coin flip, while transitions from the latter are chosen nondeterministically, perhaps in response to a stimulus from the environment.

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 must be* image-finite*: for every* $n \in S_N$ *and* $a \in Act$, $|\{s' \mid (n, a, s') \in \delta\}| < \infty$.
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.*

An NPLTS encodes the operational behavior of a system. 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, with the the system being able to 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 state of the system. Note that every successor of a probabilistic state p is a nondeterministic state. $I(A)$ returns 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. 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 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, where for readability we have suppressed $I(A)$.

NPLTSs were first introduced by Hansson [6]. Other state-machine-oriented models of probabilistic computation have natural encodings as NPLTSs, and we briefly review these here. *Markov decision processes* [4] (MDPs) contain a mixture of probabilistic and nondeterministic states, although there is no notion of action. Consequently, MDPs may be represented as NPLTSs whose action set $Act = \{a\}$ is a singleton set. *Markov chains* contain only probabilistic transitions; these may be represented as NPLTSs with $Act = \{a\}$ and with δ satisfying the property that $\delta \subseteq S_N \times Act \times S_P$ is a function of type $S_N \times Act \to S_P$. *Reactive probabilistic labeled transition systems* [9] may be encoded as NPLTSs in which $\delta \subseteq S_N \times Act \times S_P$ is a partial function over $S_N \times Act$. Nondeterministic labeled transition systems may be obtained by taking $S_P = \emptyset$.

3 A Probabilistic Semantics for the Mu-Calculus

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

3.1 NPLTS as Experiments

We treat an NPLTS as a probabilistic experiment 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 presents an example; in this case the outcome is a finite tree, but this need not be the case in general. To formalize these notions we adapt the construction given in [10] by defining the outcomes of a NPLTS as trees whose nodes are labeled by (partial) execution sequences of the NPLTS.

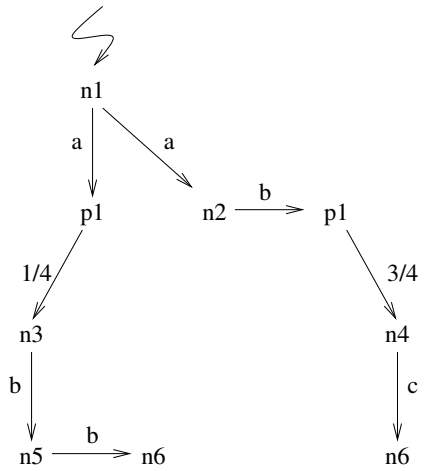

Figure 2: An outcome of the NPLTS in Figure 1.

Definition 2 *Let $L = (S_N, S_P, \delta, \mathcal{P}, I)$ be a NPLTS. A sequence $s_0 \xrightarrow{q_1} s_1 \cdots \xrightarrow{q_k} s_k$, where $0 \leq k \leq \infty$, is a* partial computation *of L if for all $0 \leq i < k$, $s_i \xrightarrow{q_{i+1}} s_{i+1}$. 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.*

The following notions 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 $\sigma' = s_0 \xrightarrow{q_1} s_1 \ldots s_k \xrightarrow{q_k} p \xrightarrow{r'} n' \ldots$, $n = n'$ and $r = r'$. (Recall $p \in S_P, n \in S_N$.)*

We also adopt the usual definitions of *prefix* and *prefix-closed*. Note that 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, and let $s \in S$. Then $\emptyset \neq T \subseteq \mathcal{C}_L(s)$ is an* observation tree rooted at *s if T is prefix-closed and observable.*

If T is an o-tree then we use $\mathsf{root}(T)$ to refer to its root. 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 $\mathcal{M}_L$ to refer to the set of all maximal o-trees of L and $\mathcal{M}_L(s) \subseteq \mathcal{M}_L$ for the maximal trees rooted at s. It is precisely the set $\mathcal{M}_L$ that we take to be the "outcomes" of the probabilistic experiment defined by L.

Intuitively, $T \stackrel{a}{\rightarrow} T'$ holds if T' is a subtree of T reachable via an edge labeled by a. This definition also reflects a key fact about NPTLS: probabilistic "transitions" are not execution steps in the same way that transitions labeled by actions are.

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

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

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 $B_T = \{T' \in \mathcal{M}_L \mid T \subseteq T'\}$.*

The measure of a basic cylindrical set B_T is defined as the proportion of all outcomes emanating from the root of B_T that have B_T as a prefix.

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

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. $\mathsf{m}_s : \mathcal{B}(s) \rightarrow [0,1]$ is then defined inductively as follows.*

$$\begin{array}{rcl} \mathsf{m}_s(B_T) & = & \mathsf{m}(B_T) \\ \mathsf{m}_s(\bigcup\limits_{i \in I} B_i) & = & \sum\limits_{i \in I} \mathsf{m}_u(B_i) \;\text{ for pairwise disjoint } B_i \\ \mathsf{m}_s(B^c) & = & 1 - \mathsf{m}_s(B) \end{array}$$

It is easy to show that m_s is a probability measure over $\mathcal{B}(s)$; thus $(\mathcal{M}_L(s), \mathcal{B}(s), \mathsf{m}_s)$ is a probability space. Set $M \subseteq \mathcal{M}_L(s)$ is *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 labeled transition systems: the semantics assigns to each formula a set of states for which the formula "holds" [8]. 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 an NPLTS L, any formula may be interpreted as sets of outcomes. Then we establish that these sets are measurable. 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 requires explanation. A tree can "satisfy" A in one of two ways. If the tree's root is a nondeterministic state satisfying A, then the tree does. On the other hand, if the tree's root is a probabilistic state, and hence "ephemeral", and the (by definition sole) transition from it leads to a nondeterministic state satisfying A, then the tree does as well. 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 [13]. 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 [8]. 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.

The next theorem establishes that for an alternation-free formula ψ and state s, the set $\Theta_{L,s}(\psi)e$ of outcomes rooted at s and satisfying ψ is measurable. This implies that $\mathsf{m}_s(\Theta_{L,s}(\psi)e)$ is well-defined; we take this quantity to be the "probability" that system L, starting 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 an alternation-free mu-calculus formula. Then $\Theta_{L,s}(\psi)e$ is measurable.*

The argument relies on the continuity μ and co-continuity of ν. The line of reasoning breaks when non-alternation-free formulas are allowed, since for example ν is not continuous. We nevertheless conjecture that that the full mu-calculus is measurable.

4 Model Checking

We now present a procedure that, given an NPLTS L, a state s in L, and an alternation-free mu-calculus formula, computes $\mathsf{m}_s(\psi)$. The approach first builds a graph from L, s and ψ, derives equations from it, and then and solves the equations. 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 labels nodes with either a *state assertion* or *a transition assertion*. Both are interpreted semantically as sets of trees.

Definition 10 (State assertion) *A* state assertion *is a pair (s, F) where $s \in S$ and $F \subseteq \Psi$ is a finite set of formulae. We define $[\![(s,F)]\!] = \mathcal{M}_L(s) \cap \Theta_L(\wedge F)$* ■

Definition 11 (Transition assertion) *A* transition assertion *is a set $\rho \subseteq S_N \times \Psi \times S$ satisfying: (a) if $(n,\psi,s),(n',\psi',s') \in \gamma$ then $n = n'$; (b) if $(n,\psi,s) \in \gamma$ then ψ has form $\langle a\rangle\psi'$ or $[a]\psi'$;and (c) if $(n,\langle a\rangle\psi,s)/(n,[a]\psi,s) \in \gamma$ then $n \xrightarrow{a} s$. We define $[\![(n,\langle a\rangle\psi,s)]\!] = [\![(n,[a]\psi,s)]\!] = \mathsf{f}_{\langle a\rangle}(\Theta_{L,n}(\psi)) \cap \mathcal{M}_L(n)$, where $\mathsf{f}_{\langle a\rangle}(M) = \{T \in \mathcal{M}_L \mid \exists T' \in M.\, T \xrightarrow{a} T'\}$, and $[\![\rho]\!] = \bigcap_{(n,\psi,s)\in\gamma}[\![(n,\psi,s)]\!]$.* ■

We now introduce the following operations on transition assertions.

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

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

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

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

■

Given state s and formula ψ, the graph is constructed inductively beginning with a node labeled by state assertion $(s, \{\psi\})$. Nodes containing transition assertions arise as a result of processing *modal state assertions*, i.e. state assertions (s, F) where all formulas in F have form $\langle a \rangle \psi$ or $[a]\psi$. Since $[\![(n, F)]\!] = \Theta_{L,n}(\bigwedge F)$ 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* $\gamma = (n, F)$*, transition assertion* ρ *is* reasonable *with respect to* γ *(written as* $\mathsf{Reasonable}(\gamma, \rho)$*) provided:*

- *For every* $(n, \psi, s) \in \rho$, $\psi \in F$ *and* $trans(\{(n, \psi, s)\}) \subseteq \delta$.
- *For every* $[a]\psi$ *in* F, $\{(n, [a]\psi, s) \mid n \xrightarrow{a} s\} \subseteq \rho$, *and for every* $\langle a \rangle \psi$ *in* F, $\{(n, \langle a \rangle \psi, s) \mid n \xrightarrow{a} s\} \cap \rho \neq \emptyset$.

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

Note that 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 $rank(\gamma, \gamma')$ 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* $\gamma = (n, F)$ *and* ρ *be such that* $\mathsf{Reasonable}(\gamma, \rho)$*. 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 now describe the construction of the graph $\Pr(L, s, \psi) = (\Gamma, \Delta, \ell)$, where $\Gamma \subseteq (S \times 2^{\Psi}) \cup 2^{S_N \times \Psi \times S}$ is the set of nodes, $\Delta \subseteq \Gamma \times (Act \cup \{\varepsilon\} \cup \{r \mid 0 < r \leq 1\}) \times \Gamma$ is the set of edges, and $\ell : \Delta \to \{+1, -1\}$ is the edge-labeling function. The sets Γ and Δ and the function ℓ are defined inductively as follows.

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

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

2. If $\gamma = (p, F) \in \Gamma$ for some $p \in S_P$ then for each transition $p \xrightarrow{r} n$ in L add $\gamma' = (n, F)$ to Γ, and (γ, r, γ') to Δ, and set $\ell(\gamma, r, \gamma') = +1$.

3. If $\gamma = (n, F) \in \Gamma$ some $n \in S_N$ and γ is a modal assertion then for each $\rho \in \mathcal{R}(\gamma)$ add ρ to Γ and $(\gamma, \varepsilon, \rho)$ to Δ, and set $\ell(\gamma, \varepsilon, \rho)$ to $rank(\gamma, \rho)$.

4. If $\gamma = (n, F) \in \Gamma$ some $n \in S_N$ and γ is not modal, then apply the first applicable rule below.

   (a) If there is a $\psi \in F$ such that $\psi \in Prop$ and $n \in I(\psi)$, or $\psi = \neg A$ some $A \in Prop$ and $n \notin I(A)$, or $\psi = [a]\psi'$ and $n \stackrel{a}{\not\rightarrow}$, then add $\gamma' = (n, F - \{\psi\})$ to Γ and $(\gamma, \varepsilon, \gamma')$ to Δ, and set $\ell(\gamma, \varepsilon, \gamma') = +1$.

   (b) If there is a $\psi \in F$ such that $\psi \in Prop$ and $n \notin I(A)$, or $\psi = \neg A \in F$ and $n \in I(A)$, or $\psi = \langle a \rangle \psi' \in F$ and $n \stackrel{a}{\not\rightarrow}$, then add $\gamma' = (n, \{false\})$ to Γ and $(\gamma, \varepsilon, \gamma')$ to Δ, and set $\ell(\gamma, \varepsilon, \gamma') = +1$.

   (c) If there is a $\psi = \psi_1 \wedge \psi_2 \in F$ then add $\gamma' = (n, (F - \{\psi\}) \cup \{\psi_1, \psi_2\})$ to Γ and $(\gamma, \varepsilon, \gamma')$ to Δ, and set $\ell(\gamma, \varepsilon, \gamma') = +1$.

   (d) If there is a $\psi = \psi_1 \vee \psi_2 \in F$ then add $\gamma_1 = (n, (F - \{\psi\}) \cup \{\psi_1\})$, $\gamma_2 = (n, (F - \{\psi\}) \cup \{\psi_2\})$, and $\gamma_3 = (n, (F - \{\psi\}) \cup \{\psi_1 \wedge \psi_2\})$ to Γ and $(\gamma, \varepsilon, \gamma_i)$ for $1 \leq i \leq 3$ to Γ, and set $\ell(\gamma, \varepsilon, \gamma_1) = \ell(\gamma, \varepsilon, \gamma_2) = +1$ and $\ell(\gamma, \varepsilon, \gamma_3) = -1$.

   (e) If there is a $\psi = \nu X.\psi' \in F$ or $\psi = \mu X.\psi' \in F$ then add $\gamma' = (s, (F - \{\psi\}) \cup \{\psi'[X \mapsto \psi]\})$ to Γ and $(\gamma, \varepsilon, \gamma')$ to Δ, and set $\ell(\gamma, \varepsilon, \gamma') = 1$.

5. If $\rho \in \Gamma$ is a transition assertion with a common first component $n \in S_N$ then for every $n \xrightarrow{a} s \in trans(\rho)$ add $\gamma' = (n, residue(\rho, n \xrightarrow{a} s))$ to Γ and (ρ, a, γ') to Δ, and set $\ell(\rho, a, \gamma') = +1$.

For any finite-state L, s and ψ $\Pr(s, \{\psi\})$ is finite due to the finiteness of the NPLTS and the fact that the all formulae arising in the construction are elements in the Fisher-Ladner closure of ψ [8]. The next lemma characterizes the relationships between adjacent nodes in $\Pr(L, s, \psi)$.

Lemma 1 *Let γ be a node in $\Pr(L, s, \psi)$. Then the following hold.*

1. *If $\gamma = (p, F)$ where $p \in S_P$, then $[\![\gamma]\!] = \bigcup_{p \xrightarrow{r} s} \{p \xrightarrow{r} T \mid T \in [\![(s, F)]\!]\}$ and $\mathsf{m}_p([\![\gamma]\!]) = \sum_{p \xrightarrow{r} n} r \times \mathsf{m}_s([\![(n, F)]\!])$*

2. *If $\gamma = (n, F)$ has successors $\gamma_1, \ldots \gamma_k$ then $[\![\gamma]\!] = \bigcup_{i=1}^{k} [\![\gamma_i]\!]$ and $\mathsf{m}_s([\![\gamma]\!]) = \sum_{i=1}^{k} \ell(\gamma, \varepsilon, \gamma_i) \times \mathsf{m}_s([\![\gamma_i]\!])$*

3. *If γ is a transition assertion with a common first component $n \in S_N$ then*

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

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

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

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:

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

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 may be other such solutions, so we must indicate how compute the ones we want. The following procedure does so.

1. Compute the strongly connected components of the graph from which the equations are constructed and topologically sort them.
2. Propagate solutions: if a solution has been computed for a variable, substitute if for all right-hand-side occurrences of 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 fall within a tolerance ε, which is a parameter to the algorithm.
   (c) Propagate these values.

This algorithm is approximation-based and therefore requires the specification of an "error tolerance" ε. However, all the functions being used are continuous, and lattice-theoretic arguments ensure that the iteration process 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 [10] 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 [3] 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

[1] AZIZ, A., SINGHAL, V., BALARIN, F., BRAYTON, R., AND SANGIOVANNI-VINCENTELLI, A. It usually works: The temporal logic of stochastic systems. In *Proc. Comp. Aided Verif.* (July 1995), Springer-Verlag.

[2] BIANCO, A., AND DE ALFARO, L. Model checking of probabilistic and mondeterministic systems. In *Found. Software Technology and Theoret. Comp. Sci.* (Dec. 1995), vol. 1026 *LNCS*, Springer-Verlag, pp. 499–513.

[3] COURCOUBETIS, C., AND YANNAKAKIS, M. Verifying temporal properties of finite-state probabilistic programs. In *Proc. 1988 IEEE Symp. on the Found. of Comp. Sci.* (1988).

[4] DERMAN, C. *Finite State Markov Decision Processes*. Academic Press, 1970.

[5] EMERSON, E. A., AND HALPERN, J. Y. "sometimes" and "not never" revisited: On branching time versus linear time temporal logic. *JACM 33*, 1 (1986), 151–178.

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

[7] HUTH, M., AND KWIATKOWSKA, M. Quantitative analysis and model checking. In *Proc. IEEE Symp. on Logic in Comp. Sci.* (Jun.–Jul. 1997), IEEE Computer Society Press, pp. 111–122.

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

[9] LARSEN, K. G., AND SKOU, A. Bisimulation through probabilistic testing. *Information and Computation 94*, 1 (1991), 1–28.

[10] NARASIMHA, M., CLEAVELAND, R., AND IYER, P. Probabilistic temporal logics via the modal mu-calculus. In *Proc. Founda. of Software Sci. and Computat. Structures* (Mar 1999), W. Thomas, Ed., ETAPS, Springer, pp. 288–305.

[11] PNUELI, A., AND ZUCK, L. D. Probabilistic verification. *Information and Computation 103*, 1 (Mar. 1993), 1–29.

[12] SEGALA, R. A compositional trace-based semantics for probabilistic automata. In *CONCUR95* (1995), pp. 324–338.

[13] TARSKI, A. A fixed point theorem and its applications. *Pacific J. Math.* (1955), 285–309.

[14] VARDI, M. Y. Automatic verification of probabilistic concurrent finite-state programs. In *IEEE Symp. on Found. of Comp. Sci.* (1985), pp. 327–338.

Rance Cleaveland is a Professor in the Department of Computer Science, SUNY at Stony Brook, Stony Brook, NY 11794-4400, USA. E-mail: rance@cs.sunysb.edu.

S. Purushothaman Iyear is an Associate Professor in the Department of Computer Science, North Carolina State University, Raleigh, NC 27695-7534, USA. E-mail: purush@csc.ncsu.edu.