

Linear-Time and May-Testing in a Probabilistic Reactive Setting[★]

Lucia Acciai, Michele Boreale, and Rocco De Nicola

Dipartimento di Sistemi e Informatica, Università degli Studi di Firenze, Italy
{lucia.acciai, michele.boreale, rocco.denicola}@unifi.it

Abstract. We consider reactive probabilistic labelled transition systems (RPLTS), a model where internal choices are refined by probabilistic choices. In this setting, we study the relationship between linear-time and may-testing semantics, where an angelic view of nondeterminism is taken. Building on the model of *d-trees* of Cleaveland et al., we first introduce a clean model of probabilistic may-testing, based on simple concepts from measure theory. In particular, we define a probability space where statements of the form “*p may pass test o*” naturally correspond to measurable events. We then obtain an observer-independent characterization of the may-testing preorder, based on comparing the probability of *sets* of traces, rather than of *individual* traces. This entails that may-testing is strictly finer than linear-time semantics. Next, we characterize the may-testing preorder in terms of the probability of satisfying safety properties, expressed as languages of infinite *trees* rather than traces. We then identify a significative subclass of RPLTS where linear and may-testing semantics do coincide: these are the *separated* RPLTS, where actions are partitioned into probabilistic and nondeterministic ones, and at each state only one type is available.

Keywords: probabilistic transition systems, linear time, testing equivalence, safety.

1 Introduction

In a classical nondeterministic setting, it is well-known that trace equivalence is totally insensitive to points of choice in time. This makes trace equivalence a *linear*-time semantics, as opposed to the various *branching*-time semantics of van Glabbeek’s spectrum [13], ranging from bisimilarity to failure equivalence. The insensitiveness to points of choice makes linear time the ideal framework when one is interested in analyzing properties that can be expressed as prefix-closed sets of traces, like Safety.

In this context, the *testing equivalence* approach [9] is conceptually important, for it provides a clean observational justification of linear time. Indeed, in the setting of ccs and labelled transition systems, trace equivalence does coincide with *may*-testing equivalence, which deems two processes equivalent when no system (observer) running in parallel may possibly note any difference between them (*must*-testing, on the other


[★] Corresponding author: Lucia Acciai, DSI - Università di Firenze, Viale Morgagni 65, 50134 Firenze. Work partially supported by the EU project ASCENS under the FET open initiative in FP7.

hand, gives rise to failure semantics, see [8]). However expected, this coincidence result should not be taken for granted in other settings. For example, the result breaks down as soon as one moves from a synchronous to an asynchronous communication mechanism (see [3]).

In this paper, we study linear time and may-testing in a setting where internal choices are refined by probabilistic ones, the *reactive probabilistic labelled transition systems* (RPLTS for short, reviewed in Section 3) as described e.g. in [5]. RPLTS's are equivalent to the *Markov Decision Processes* used in probabilistic verification [2]. The reason for choosing this model, which only features external nondeterminism, is that the probabilistic version of linear-time semantics would not make sense in the presence of internal nondeterminism (see Section 3).

The motivation of our study is twofold. First, we are interested in formulating a *clean* probabilistic adaption of the original proposal of testing semantics [9]. This will allow us to check if, or under what assumptions, may-testing still provides an observational justification for (probabilistic) linear-time semantics. Second, with our model at hand, we hope to shed some light in some issues raised by existing probabilistic testing theories, particularly the issue of *overestimation* of probabilities (see e.g. [11] and references therein). These motivations are further discussed below.

The need for a clean model of probabilistic testing can be illustrated by the following example. Consider the two processes s and s' on the right. The two processes and the environment can synchronize over the actions $init$, l , r , ζ and \checkmark . In both s and s' , after an initial synchronization on $init$, the environment is offered a choice between l and r , then either \checkmark or ζ may be offered, also depending on the previous choice (l or r) of the environment. The interpretation of probabilistic choice is that a fair coin is internally tossed to decide which branch will be made available to the environment. The probability of any individual trace – that is, the probability that any given sequence of synchronizations becomes available to the environment – is the same for both s and s' . Hence the probabilistic linear-time semantics deems s and s' as equivalent. Indeed, following Georgievska and Andova [11], one can argue that not equating s and s' would imply overestimating the probability of success for an external observer, allowing him to observe some points of choice. On the other hand, consider the case of s and s' being two candidate implementations of a safety-critical system; here ζ represents some catastrophic event, like a system crash. In this case, one is more interested in the chance that, by resolving the internal choices, the mere possibility of ζ is ruled out, whatever the behaviour of the environment and of the scheduler. In other words, one is interested in the probability that *none* of the dangerous traces in the set $\{init \cdot l \cdot \zeta, init \cdot r \cdot \zeta\}$ becomes available. Now, s assigns to this event probability 0, while in the case of s' , the choice of the left branch of l and of the right branch of r , an event that happens with probability $\frac{1}{4}$, will rule out the possibility of ζ . In this sense, s' might be considered as safer than, and not equivalent to, s - despite the fact that this implies observing a point of choice.



An assessment of this and similar issues should rest on a conceptually clean model of testing. Over the years, many models of probabilistic testing have been put forward by a number of authors [17,21,19], up to and including the recent work by Deng et al. [6,7]. A common paradigm is compose-and-schedule, by which the nondeterminism resulting from the synchronized composition of the process and the observer¹ is resolved by employing - implicitly or explicitly - *schedulers* that make the system fully probabilistic. We see two difficulties with schedulers. First, schedulers can look at the states of the synchronized system, including “ghost” states that result from purely probabilistic choices. This may lead to an unrealistic observation power - the issue of overestimation discussed e.g. by [11]. Some models rectify this by hiding the random choices from the scheduler [4,11]. But then it may become problematic to decide which points of choice should remain observable and which should not (see the example above). Second, the outcome of testing a process with an observer is a *range* of success probabilities, one for each possible scheduler. Comparing two systems on the basis of these ranges is in general awkward. Say one deems system *A* safer than system *B* if for each scheduler of *A* there is a scheduler of *B* that will lead to crash with a greater probability (see e.g. [7]). The relation “*A* safer than *B*” thus established may be of no use in a real-world context, where both the behaviour of the environment and the actual scheduling policy are unpredictable to a large extent. This is of course a drawback when analyzing safety-critical systems. To sum up, in these approaches taking schedulers explicitly into account makes the very concept of *passing a test* awkward, and somehow spoils the clarity of the original testing proposal [9].

In this paper we face these issues from a different perspective and put forward a model that abstracts away from schedulers. The basic idea is that one should first resolve probabilistic choices, then treat the resulting nondeterministic system angelically (if an event may happen, it will happen). Informally, resolving the probabilistic choices out of a process *p* and of an observer *o* yields a pair of nondeterministic systems, *T* and *U*, with certain associated probabilities, $\Pr(T)$ and $\Pr(U)$. Here, *T* may or may not satisfy *U* in a traditional sense. Approximately, the probability that *p* may pass test *o* could then be expressed as a sum

$$\Pr(p \text{ may pass } o) = \sum_{T,U:T \text{ may pass } U} \Pr(T) \cdot \Pr(U). \quad (1)$$

We formalize this intuition building on simple concepts from measure theory (reviewed in Section 2) and on the model of *d-trees* of Cleaveland et al. [5] (reviewed in Section 4). In particular, we introduce a probability space where the statements “*p* may pass *o*” naturally correspond to measurable events. In general, the sum (1) becomes a proper integral in this space. Going back to the example above, *s* and *s'* are distinguished in this model by the observer $o = \text{init}.(l.\zeta.\omega + r.\zeta.\omega)$ (here ω is the success action): indeed, the probability that *s* may pass *o* is 1, while it is $\frac{3}{4}$ for *s'*.

With this model at hand, we investigate the relationships existing among may-testing, linear-time semantics and safety properties. In summary, we offer the following contributions:

¹ Of course, the nondeterminism arising from this composition is always of *internal* type, despite the fact that the system and the process alone may only feature external nondeterminism.

- a clean model of probabilistic may-testing for RPLTS (Subsection 4.1);
- an observer-independent characterization of the may-testing preorder, based on comparing the probability of *sets* of traces, rather than of *individual* traces (Subsection 4.2);
- a comparison of may testing with both linear-time and tree-unfolding semantics (Subsection 4.3);
- a characterization of the may-testing preorder in terms of safety properties, expressed as sets of *infinite trees* rather than traces (Section 5).
- sufficient conditions on RPLTS's and observers guaranteeing that linear and may-testing semantics do coincide. This leads to the class of *separated* RPLTS, where probabilistic and nondeterministic transitions do not mix up (Section 6).

We end the paper with a few considerations on further and related work (Section 7). For lack of space most proofs are omitted in this short version; they can be found in the full version available online [1].

2 Background on Measure Theory

We recall some notions from elementary measure theory. A classical reference is [14]. Let X be any nonempty set. A *sigma-algebra*, or *measurable space*, on X is a pair (X, \mathcal{A}) such that $\emptyset \neq \mathcal{A} \subseteq 2^X$ is closed under countable unions and complementation. A *measure* over (X, \mathcal{A}) is a function $\mu : \mathcal{A} \rightarrow \mathbb{R}^+ \cup \{\infty\}$ satisfying additivity under countable disjoint unions and such that $\mu(\emptyset) = 0$. It is a *probability measure* if $\mu(X) = 1$. The triple (X, \mathcal{A}, μ) is called *measure space*; if μ is a probability measure, it is also called a *probability space*. Let (X, \mathcal{A}, μ_1) and (Y, \mathcal{B}, μ_2) be two measure spaces. The product sigma-algebra $\mathcal{A} \times \mathcal{B}$ is defined to be the sigma-algebra on the cartesian product $X \times Y$ generated by the subsets of the form $A \times B$, with $A \in \mathcal{A}$ and $B \in \mathcal{B}$.

Given two sigma-finite [14] measure spaces (X, \mathcal{A}, μ_1) and (Y, \mathcal{B}, μ_2) , the product measure $\mu^{\mathcal{A} \times \mathcal{B}}$ is defined to be the unique measure on the measurable space $(X \times Y, \mathcal{A} \times \mathcal{B})$ satisfying the following condition

$$\mu^{\mathcal{A} \times \mathcal{B}}(A \times B) = \mu_1(A) \cdot \mu_2(B) \quad (2)$$

for all $A \in \mathcal{A}$ and $B \in \mathcal{B}$. If μ_1 and μ_2 are probability measures, so is their product $\mu^{\mathcal{A} \times \mathcal{B}}$, hence in this case $(X \times Y, \mathcal{A} \times \mathcal{B}, \mu^{\mathcal{A} \times \mathcal{B}})$ forms a probability space.

3 Reactive Probabilistic Labeled Transition Systems

This section introduces the object of our study, Reactive Probabilistic Labeled Transition Systems (RPLTS for short) and the linear-time and tree-unfolding semantics. The relationship between the linear-time and tree-unfolding preorders and the may-testing preorder will be investigated in the next section.

3.1 RPLTS

Let us fix a nonempty set Act of actions, ranged over by a, b, \dots . We will let w, v, \dots range over Act^* . A *Reactive Probabilistic Labeled Transition System* [5,12,18] is basically a finite-branching probabilistic LTS's over the set Act . Labels on transitions record the interactions the system may engage in with the environment: at each state, any given action may or may not be available for interaction. Internal nondeterminism is refined by probabilistic choices: if available, a given action can lead to different states depending on probabilities attached to transitions.

Definition 1 (RPLTS). A reactive probabilistic labeled transition system L is a triple (S, δ, P) , such that:

- S is an at most countable set of states;
- $\delta \subseteq S \times Act \times S$ is a transition relation such that for each $s \in S$ there exists a finite number of transitions of the form $(s, \cdot, \cdot) \in \delta$ (i.e. δ is finitely-branching);
- $P : \delta \rightarrow (0, 1]$ is a transition probability distribution such that for each $s \in S$ and $a \in Act$: $\sum_{s' : (s, a, s') \in \delta} P(s, a, s') \in \{0, 1\}$;

A RPLTS can be depicted as a graph, as shown on the right. Let us now introduce some terminology. Let $L = (S, \delta, P)$ be a RPLTS. We will often write $s \xrightarrow{a} s'$ to mean that $(s, a, s') \in \delta$, if the underlying L is clear from the context.

A *computation* of L is any sequence σ of the form $s_0 a_1 s_1 a_2 \dots a_n s_n \in S \cdot (Act \cdot S)^*$, where $n \geq 0$ and for each $0 \leq i < n$ it holds that $s_i \xrightarrow{a_{i+1}} s_{i+1}$. We will denote by $\text{fst}(\sigma)$ and $\text{lst}(\sigma)$, respectively, the initial and the final state of σ and by $\lambda(\sigma)$ the sequence of labels in σ , that is $\lambda(\sigma) = a_1 a_2 \dots a_n$. We define the *weight* of σ as $\text{wt}(\sigma) \triangleq \prod_{i=0}^{n-1} P(s_i, a_{i+1}, s_{i+1})$. Let us fix a RPLTS $L = (S, \delta, P)$ and any state s of S . In what follows, we will denote by \mathcal{C}^L the set of all computations over L , and by \mathcal{C}_s^L the set of all computations σ over L such that $\text{fst}(\sigma) = s$.


A computation σ' is said to be a *prefix* of σ if σ' is a prefix of σ as a string. A set of computations $D \subseteq \mathcal{C}^L$ is said to be *prefix-closed* if for every $\sigma \in D$ and σ' prefix of σ , $\sigma' \in D$. A set of computations $D \subseteq \mathcal{C}^L$ is said to be *deterministic* if whenever $\sigma, \sigma' \in D$, with $\sigma = \sigma'' a s$ and $\sigma' = \sigma'' a' s'$, then either $a \neq a'$ or $s = s'$.

Definition 2 (d-trees). Let L be a RPLTS. Then $\emptyset \neq D \subseteq \mathcal{C}^L$ is a d-tree if the following hold:

- (1) there is an $s \in S$ such that $D \subseteq \mathcal{C}_s^L$;
- (2) D is prefix-closed;
- (3) D is deterministic.

We say that a d-tree D of L is *rooted at* s if $D \subseteq \mathcal{C}_s^L$ and let \mathcal{T}^L and \mathcal{T}_s^L denote, respectively all d-trees of L and all d-trees of L rooted at s . We will write \mathcal{F}^L for the set of all finite d-trees of L and \mathcal{F}_s^L for the subset of those rooted at s . Finally, we define the weight of a $D \in \mathcal{F}_s^L$ as

$$\text{wt}(D) \triangleq \prod_{\sigma \in D} P(\text{lst}(\sigma), a, s).$$



Note that if D is given as the prefix-closure of some computation σ , then $\text{wt}(D) = \text{wt}(\sigma)$. Finally, given any d-tree $D \subseteq \mathbb{C}_s^L$, we set $\lambda(D) \triangleq \{\lambda(\sigma) \mid \sigma \in D\}$ and use $D \xrightarrow{w}$ as an abbreviation of $w \in \lambda(D)$.

3.2 Linear-Time Semantics of RPLTS

Definition 3 (probabilistic linear-time preorder). *Let L be a RPLTS. For any state s , the function $f_s^L : \text{Act}^* \rightarrow [0, 1]$ is defined thus*

$$\text{for each } w \in \text{Act}^*, \quad f_s^L(w) \triangleq \sum_{\sigma \in \mathbb{C}_s^L : \lambda(\sigma)=w} \text{wt}(\sigma). \quad (3)$$

For any two states s and s' , we write $s \leq_{\text{lin}} s'$ if and only if for each $w \in \text{Act}^$, $f_s^L(w) \leq f_{s'}^L(w)$.*

Note that the sum in (3) is finite, as L is finitely branching. Functions of the type $\text{Act}^* \rightarrow \mathbb{K}$ are classically known as *formal power series* in Automata Theory: they represent a generalization of the set-theoretic notion of language to a setting where weights of transitions are not just 0/1 (absence/presence), but elements of a semiring, \mathbb{K} (in our case, the reals). In our scenario, a natural interpretation of “ $f_s^L(w) = p$ ” is that, starting at s , with probability p a sequence of synchronizations along w will be available to an observer. Note that when applied to general PLTS², also featuring internal nondeterminism, this definition would not make sense: indeed, one might end up having $f_s^L(w) > 1$.

3.3 Tree-Unfolding Semantics of RPLTS

Some terminology on trees is in order. A *tree* θ over Act is a nonempty, prefix-closed subset of Act^* . In what follows, we shall call \mathfrak{T}^f the set of finite trees over Act^* and use the letter t to range over finite trees.

Definition 4 (probabilistic tree-unfolding preorder). *Let L be a RPLTS. For any state s , the function $\varphi_s^L : \mathfrak{T}^f \rightarrow [0, 1]$ is defined thus*

$$\text{for each } t \in \mathfrak{T}^f, \quad \varphi_s^L(t) \triangleq \sum_{D \subseteq \mathbb{C}_s^L : \lambda(D)=t} \text{wt}(D). \quad (4)$$

For any two states s and s' , we write $s \leq_{\text{tree}} s'$ if and only if for each $t \in \mathfrak{T}^f$, $\varphi_s^L(t) \leq \varphi_{s'}^L(t)$.

Note that the sum in (4) is finite, as L is finitely branching and t is finite. Functions of type $\mathfrak{T}^f \rightarrow \mathbb{K}$ are known as *formal tree series* in Automata Theory (see e.g. [10]) and represent a generalization of formal power series to trees.

By Definition 3 and 4 it follows that the tree-unfolding preorder is included in the linear-time preorder. The example in the Introduction witnesses the fact that this inclusion is strict: the linear-time semantics deems s and s' as equivalent, while the tree-unfolding semantics does not. Indeed, $\varphi_s^L(t) > \varphi_{s'}^L(t)$ and $\varphi_s^L(t') < \varphi_{s'}^L(t')$, with $t = \{\epsilon, \text{init}, \text{init}\cdot l, \text{init}\cdot r, \text{init}\cdot l\cdot\checkmark, \text{init}\cdot r\cdot\checkmark\}$ and $t' = \{\epsilon, \text{init}, \text{init}\cdot l, \text{init}\cdot r, \text{init}\cdot l\cdot\checkmark, \text{init}\cdot r\cdot\checkmark\}$. We sum the above discussion in the following:

² These can be obtained from Definition 1 by replacing the condition of the third item with just $\sum_{s' : (s, a, s') \in \delta} P(s, a, s') \in \mathbb{N}$.

Proposition 1. *The preorder \leq_{tree} is strictly included in \leq_{lin} .*


4 May Testing for RPLTS

In the first part, we review the probabilistic model of d-trees introduced by Cleaveland et al. [5]. In the second part, we introduce the testing scenario and the definition of may-testing preorder. We then establish a trace-based, observer-independent characterization of this preorder.

4.1 The Sigma-Algebra of d-trees

The material in this subsection is borrowed from [5]. Some additional terminology is in order. In what follows, we fix a generic RPLTS $L = (S, \delta, P)$. Given any $D, D' \subseteq \mathcal{C}^L$ we say that D' is a prefix of D if $D' \subseteq D$. A d-tree is said *maximal* if it is not prefix of any other d-tree; we write \mathcal{M}^L and \mathcal{M}_s^L , respectively, for the set of all maximal d-trees of L and of all maximal d-trees of L rooted at s . In what follows, we let T and U range over maximal d-trees.

Example 1. Consider the RPLTS L depicted in Section 3; the d-trees below belong to $\mathcal{F}_{s_0}^L$.



In the picture, each path from the root to a node of the tree - either leaf or internal node - represents a computation in the corresponding d-tree. Notice also that D_1 is a prefix of D_0 , therefore it does not belong to \mathcal{M}_s^L , while D_0 and D_2 do.

Following [5], we consider the d-trees in \mathcal{M}^L as the possible “outcomes” of observing L , and according to this intuition, define a probability space over \mathcal{M}_s^L , for a given state s of L . The construction of this space is based on the concept of “basic cylindrical sets” of maximal d-trees: subsets of \mathcal{M}_s^L containing all d-trees sharing a given finite prefix. The measure of each cylinder is defined, as expected, as the product of the probabilities associated to edges of the common prefix. Formal definitions of basic cylindrical sets and of their measure are given below.

Definition 5 (basic cylindrical set). *Let $L = (S, \delta, P)$ be a RPLTS, let $s \in S$ and $D \in \mathcal{F}_s^L$. The basic cylindrical set (with common prefix D) $B_D \subseteq \mathcal{M}_s^L$ is defined as: $B_D \triangleq \{T \in \mathcal{M}_s^L \mid D \subseteq T\}$.*

We let \mathcal{B}_s be the sigma-algebra generated by the collection of basic cylindrical sets B_D , for $D \in \mathcal{F}_s^L$. \mathcal{B}_s is obviously sigma-finite. We let $\mu_s^L : \mathcal{B}_s \rightarrow [0, 1]$ be the unique measure satisfying $\mu_s^L(B_D) = \text{wt}(D)$ for each $D \in \mathcal{F}_s^L$.

For any s , μ_s^L is a probability measure over \mathcal{B}_s , therefore $(\mathcal{M}_s^L, \mathcal{B}_s, \mu_s^L)$ is a probability space. In the following, we will omit the superscript L from μ_s^L when the underlying RPLTS is clear from the context.

4.2 The May-Testing Preorder

Let us fix a generic RPLTS L . An *observer* O is a RPLTS over the set of actions $Act \cup \{\omega\}$, where $\omega \notin Act$ is a distinct *success* action. For any state o of O and $H \in \mathcal{T}_o^O$, we let $\Omega(H) \triangleq \{w \in Act^* \cdot \{\omega\} \mid H \xrightarrow{w}\}$. The set of all possible sequences of actions leading o to success is then $\Omega(o) \triangleq \bigcup_{H \in \mathcal{T}_o^O} \Omega(H)$. A state o of O is said to be *finite* if \mathcal{C}_o^O is a finite set. In other words, o is finite if the RPLTS reachable from o is finite-state and acyclic. In the following we write W° for the set $\{w \mid w\omega \in W\}$.

Definition 6. Let s and o be states of L and of O , respectively. For any $D \in \mathcal{T}_s^L$ and $H \in \mathcal{T}_o^O$, we say that D may H if there is $w \in Act^*$ such that $D \xrightarrow{w}$ and $H \xrightarrow{w\omega}$. The set of maximal d-trees of L rooted at s that satisfy o is defined as

$$\text{sat}(s, o) \triangleq \{(T, U) \in \mathcal{M}_s^L \times \mathcal{M}_o^O \mid T \text{ may } U\}.$$

Before introducing the may-testing preorder, we have to fulfill one obligation: proving that $\text{sat}(s, o)$ is measurable in an appropriate sigma-algebra. To this purpose, we first generalize the concept of maximal d-tree. Informally, a W -maximal d-tree, with $W \subseteq Act^*$, is a d-tree D such that $D \xrightarrow{w}$ for at least one $w \in W$. Moreover, D is non redundant with respect to W , in the sense that it cannot be extended (resp. pruned) to match more (resp. the same) elements of W . These conditions, plus a requirement of local maximality on nodes, ensure that distinct W -maximal d-trees generate disjoint basic cylindrical sets.

Definition 7 (W-maximal d-tree). Let L be a RPLTS, s a state and let $W \subseteq Act^*$. $D \in \mathcal{T}_s^L$ is said to be locally-maximal if whenever $\sigma a s_1 \in D$ and $\sigma b s_2 \in \mathcal{C}_s^L$ then there is s_3 s.t. $\sigma b s_3 \in D$. $D \in \mathcal{T}_s^L$ is said to be W -maximal if it is locally-maximal and satisfies the following conditions (the inclusions below are strict):

1. $\lambda(D) \cap W \neq \emptyset$;
2. for each locally-maximal $D' \in \mathcal{T}_s^L$, $D \subset D'$ implies $\lambda(D') \cap W = \lambda(D) \cap W$;
3. for each locally-maximal $D' \in \mathcal{T}_s^L$, $D' \subset D$ implies $\lambda(D') \cap W \subset \lambda(D) \cap W$.

This definition is extended to observers by letting W range over subsets of $Act^* \cdot \{\omega\}$.

It is worthwhile to note that the Act^* -maximal d-trees (rooted at s) are exactly the maximal d-trees (rooted at s).

Example 2. Consider again the RPLTS in Section 3 and the d-trees D_0, D_1, D_2 and D_3 from Example 1. Let $W = \{ade, ad\}$. Then:

- D_1 is not locally-maximal, hence not W -maximal: it does not contain the transition $s_2 \xrightarrow{e} s_6$;
- D_3 is not W -maximal: the transition $s_8 \xrightarrow{d} s_{10}$ is unnecessary;
- D_0 and D_2 are W -maximal.

The following key result motivates the introduction of W -maximal d-trees. It follows from the definition of basic cylindrical set, locally- and W -maximal d-tree.

Lemma 1. Let $D, D' \in \mathcal{T}_s^L$ be W -maximal d -trees, for some $W \subseteq \text{Act}^*$. If $D \neq D'$ then $B_D \cap B_{D'} = \emptyset$

We come to show that $\text{sat}(s, o)$ is measurable.

Proposition 2. The set $\text{sat}(s, o)$ is measurable in the product sigma-algebra $\mathcal{B}_s \times \mathcal{B}_o$. Moreover, if o is finite then $\text{sat}(s, o) = \biguplus_{(D,H):} \begin{array}{l} H \in \mathcal{F}_o^O \text{ } \Omega(o)\text{-maximal} \\ D \in \mathcal{F}_s^L \text{ } \Omega(H)^\circ\text{-maximal} \end{array} B_D \times B_H$.

Consider the probability spaces $(\mathcal{M}_s^L, \mathcal{B}_s, \mu_s)$ and $(\mathcal{M}_o^O, \mathcal{B}_o, \mu_o)$ and let $\mu_{(s,o)}$ denote the product probability measure over $\mathcal{B}_s \times \mathcal{B}_o$. As a corollary of Proposition 2, of the definition of product sigma-algebras and product measures (2), we get the following.

Corollary 1. For a finite o , $\mu_{(s,o)}(\text{sat}(s, o)) = \sum_{(D,H):} \begin{array}{l} H \in \mathcal{F}_o^O \text{ } \Omega(o)\text{-maximal} \\ D \in \mathcal{F}_s^L \text{ } \Omega(H)^\circ\text{-maximal} \end{array} \text{wt}(D) \cdot \text{wt}(H)$.

The classical definition of may testing preorder [9] is extended to the present probabilistic setting by taking into account the probability that two states satisfy any given observer. Note that the preorder thus defined only relates states of the same RPLTS. In case one wants to relate states belonging to two different RPLTS's, or even relate two rooted RPLTS's, one may work with the disjoint union of the two RPLTS's.

Definition 8 (may testing preorder). Let $L = (S, \delta, P)$ be a RPLTS, let $s, s' \in S$, and O be a set of observers. We define $s \sqsubseteq^{L,O} s'$ if and only if for any observer $O \in O$ and any state o in O , it holds that $\mu_{(s,o)}(\text{sat}(s, o)) \leq \mu_{(s',o)}(\text{sat}(s', o))$.

When O is the whole class of observers with actions in $\text{Act} \cup \{\omega\}$, we abbreviate $s \sqsubseteq^{L,O} s'$ just as $s \sqsubseteq^L s'$, and call this just the *may-testing* preorder. The superscript L will be omitted when clear from the context.

We can now define a trace-based, observer-independent characterization of the may-testing preorder. Let us fix a generic RPLTS $L = (S, \delta, P)$ and take $s \in S$ and any $W \subseteq \text{Act}^*$. We define $(s \xRightarrow{W}) \triangleq \{T \in \mathcal{M}_s^L \mid T \xrightarrow{w} \text{ for some } w \in W\}$ and let $(s \xRightarrow{W})$ stand for $(s \xRightarrow{W})$ with $W = \{w\}$.

Theorem 1 (observer-independent characterization). For each s and s' states of L , $s \sqsubseteq s'$ if and only if for every $W \subseteq_{\text{fin}} \text{Act}^*$, one has $\mu_s(s \xRightarrow{W}) \leq \mu_{s'}(s' \xRightarrow{W})$.

Proof: In the proof we will use the following facts (proofs can be found in [1, Appendix A]):

- (a) The set $(s \xRightarrow{W})$ is measurable in \mathcal{B}_s . In particular, if W is finite, one has $(s \xRightarrow{W}) = \biguplus_{D \in \mathcal{F}_s^L: D \text{ is } W\text{-maximal}} B_D$.
- (b) Let L be a RPLTS, O be an observer and s, o be states of L and O , respectively. For each $U \in \mathcal{M}_o^O$ define $E_{s,U} \triangleq \{T \in \mathcal{M}_s^L \mid (T, U) \in \text{sat}(s, o)\}$. Then $\mu_{(s,o)}(\text{sat}(s, o)) = \int_{\mathcal{M}_o^O} \mu_s(E_{s,U}) d\mu_o(U)$.

Assume $s \sqsubseteq s'$. Fix any $W \subseteq_{\text{fin}} \text{Act}^*$. One can build a *deterministic* observer O_W such that for some finite state o in O_W , one has $\Omega(o) = W \cdot \{\omega\}$. Since O_W is deterministic, one has that $\mathcal{M}_o^{O_W}$ consists of a single d-tree, say H , which is also the unique $W \cdot \{\omega\}$ -maximal d-tree. Moreover, $\text{wt}(H) = 1$ by definition. Now we have

$$\begin{aligned} \mu_{(s,o)}(\text{sat}(s, o)) &= \sum_{D \in \mathcal{F}_s^L, D \text{ } W\text{-maximal}} \text{wt}(D) \cdot \text{wt}(H) \text{ (by Corollary 1)} \\ &= \sum_{D \in \mathcal{F}_s^L, D \text{ } W\text{-maximal}} \text{wt}(D) \text{ (by } \text{wt}(H) = 1) \\ &= \mu_s(\bigoplus_{D \in \mathcal{F}_s^L, D \text{ } W\text{-maximal}} B_D) \text{ (by } \mu_s(B_D) = \text{wt}(D) \text{ and additivity)} \\ &= \mu_s(s \xRightarrow{W}) \text{ (by (a)).} \end{aligned}$$

Finally, $\mu_{(s,o)}(\text{sat}(s, o)) \leq \mu_{(s',o)}(\text{sat}(s', o))$ implies $\mu_s(s \xRightarrow{W}) \leq \mu_{s'}(s' \xRightarrow{W})$.

Assume now that for every $W \subseteq_{\text{fin}} \text{Act}^*$ one has $\mu_s(s \xRightarrow{W}) \leq \mu_{s'}(s' \xRightarrow{W})$. Take any observer O and any state o of O . For every $U \in \mathcal{M}_o^O$, let $V = \Omega(U)^\circ \subseteq \text{Act}^*$. The – possibly infinite – set V can be written as $V = \bigcup_{i \geq 0} V_i$, where each V_i is the subset of V containing sequences of length $\leq i$. By the properties of measures, for any r , $\mu_r(r \xRightarrow{V}) = \lim_{i \rightarrow \infty} \mu_r(r \xRightarrow{V_i})$. Since for each i , $\mu_s(s \xRightarrow{V_i}) \leq \mu_{s'}(s' \xRightarrow{V_i})$, on the limit we get $\mu_s(E_{s,U}) = \mu_s(\xRightarrow{V}) \leq \mu_{s'}(s' \xRightarrow{V}) = \mu_{s'}(E_{s',U})$. Therefore, by integrating the two functions $U \mapsto \mu_s(E_{s,U})$ and $U \mapsto \mu_{s'}(E_{s',U})$ over \mathcal{M}_o^O , it follows that

$$\int_{\mathcal{M}_o^O} \mu_s(E_{s,U}) d\mu_o(U) \leq \int_{\mathcal{M}_o^O} \mu_{s'}(E_{s',U}) d\mu_o(U).$$

This is equivalent to $\mu_{(s,o)}(\text{sat}(s, o)) \leq \mu_{(s',o)}(\text{sat}(s', o))$, by (b). Since O and o are arbitrary, it follows that $s \sqsubseteq s'$. \square

4.3 On the Relationship Among \sqsubseteq , \leq_{lin} and \leq_{tree}

We can finally make precise the relationship between the may-testing preorder \sqsubseteq , the linear and tree-unfolding preorder, \leq_{lin} and \leq_{tree} respectively. In Proposition 1 we have already shown that \leq_{tree} is strictly included in \leq_{lin} . It remains to establish a relationship between \sqsubseteq and \leq_{lin} and between \sqsubseteq and \leq_{tree} . We start by considering the first pair of preorders and by introducing a simple result that is at the basis of Theorem 2 below.

Lemma 2. *For each $s \in S$ and $w \in \text{Act}^*$, $\mu_s(s \xRightarrow{w}) = f_s(w)$.*

Theorem 2. *The preorder \sqsubseteq is strictly included in \leq_{lin} .*

Remark 1 (on canonical observers). The proof of the above Theorem 1 shows that the class of finite, deterministic and *non*-probabilistic observers of the form O_W ($W \subseteq_{\text{fin}} \text{Act}^*$) is powerful enough to induce the may testing preorder \sqsubseteq .

On the other hand, we also note that “linear” observers, i.e. observers composed by a single successful sequence, that are sufficient in the classical case, are not sufficient here. It is quite easy to see that they induce the preorder \leq_{lin} .

Consider now the RPLTS L depicted in Fig. 1. It witnesses the fact that neither $\sqsubseteq \subseteq \leq_{\text{tree}}$ nor $\leq_{\text{tree}} \subseteq \sqsubseteq$. As a matter of fact, Theorem 1 guarantees that $s \not\sqsubseteq s'$, indeed $\mu_{s'}^L(s' \xrightarrow{\{ab,ac\}}) = 0.6 < \mu_s^L(s \xrightarrow{\{ab,ac\}}) = 0.7$, while, for what concerns





Fig. 1. $s \not\leq s'$ and $s' \not\leq_{\text{tree}} s$.

the tree-unfolding preorder, we get the opposite: $s' \not\leq_{\text{tree}} s$. Indeed, $\varphi_s^L(t) = 0.1 < \varphi_{s'}^L(t) = 0.2$, where t is the tree represented by $\{\epsilon, a, ab, ac\}$.

To conclude, we pictorially represent on the right the inclusion relationships among the three preorders thus established.



5 May-Testing and the Safety Properties of Infinite Trees

A more satisfactory understanding of a behavioural relation can be obtained by looking at it in terms of the class of properties satisfied by equivalent processes. A famous example is the Hennessy-Milner theorem [16], asserting that two processes are bisimilar exactly when they satisfy the same formulae of the HM logic. Another example, in a probabilistic setting, is the characterization of probabilistic bisimulation in the work of Larsen and Skou [18]. In our case, the alternative characterization in terms of sets of traces obtained in the previous section suggests looking at properties of trees. In fact, we shall characterize the may testing preorder in terms of the probability, for two given states, of satisfying safety properties of trees.

Some additional terminology on strings and trees is in order. Recall that a (possibly infinite) tree θ is a prefix-closed subset of Act^* . Let us indicate by $<$ the usual prefix partial order on strings. The set of *leaves* of θ , denoted by $\text{leaves}(\theta)$, is the set of strings in θ that are $<$ -maximal in θ . We say a tree is *maximal* if $\text{leaves}(\theta) = \emptyset$; note that a maximal tree is necessarily infinite. We call \mathfrak{T} the set of maximal trees. In what follows, we shall use the letter τ to range over \mathfrak{T} . There is a natural partial ordering on trees given by the following

$$\theta \leq \theta' \text{ iff } \theta \subseteq \theta' \text{ and whenever } w \in \theta' \setminus \theta \text{ then there is } u \in \text{leaves}(\theta) \text{ s.t. } u < w.$$

What this means is that θ' can be obtained from θ by expanding into trees the leaves of θ . If $\theta \leq \theta'$ we also say θ is a prefix of θ' . Let us call Θ the sigma-algebra of maximal trees generated by the basic cylindrical sets C_t , where t ranges over all finite trees and $C_t \triangleq \{\tau \mid t \leq \tau\}$.

Let us now fix a RPLTS L . We shall assume that L has *no dead state*, that is, for each state s there is always at least one transition from s . This assumption allows for a simpler treatment in the following, but is not really restrictive: any RPLTS can be turned into one with no dead states by adding, where necessary, dummy self-loops labelled by a distinct action. For any state s in L , recall that \mathcal{B}_s is the sigma-algebra of maximal d-trees on L rooted at s (Section 4.1). Note that the label-extracting function $\lambda : \mathcal{B}_s \rightarrow \Theta$ maps each $T \in \mathcal{B}_s$ into a maximal tree $\tau = \lambda(T) \in \Theta$. Also note that whenever C_t

and $C_{t'}$ are disjoint, so are $\lambda^{-1}(C_t)$ and $\lambda^{-1}(C_{t'})$. As a consequence, $\lambda^{-1}(\bigsqcup_{t \in I} C_t) = \bigsqcup_{t \in I} \lambda^{-1}(C_t)$. Moreover, $\lambda^{-1}(C_t^c) = (\lambda^{-1}(C_t))^c$. Another property of λ we shall rely upon is the following:

Lemma 3. *For any t , $\lambda^{-1}(C_t)$ is measurable in \mathcal{B}_s .*

The previous properties of λ allow us to define measures on Θ as follows.

Definition 9. *Let s be a state of L . The measure ν_s on Θ is defined by setting for the basic cylindrical sets $\nu_s(C_t) \triangleq \mu_s(\lambda^{-1}(C_t))$.*

With the above definitions, $(\mathfrak{T}, \Theta, \nu_s)$ is a probability space, for each s in L . The following lemma is a consequence of the additivity of measures μ_s and ν_s and of the fact that λ^{-1} preserves disjointness.

Lemma 4. *Let $R = \bigsqcup_{t \in I} C_t$, for some index set I . Then $\nu_s(R) = \mu_s(\lambda^{-1}(R))$.*

The elements of Θ we are interested in are the *safety* properties of the form $\text{Safe}_W = \{\tau \mid \tau \cap W = \emptyset\}$, where W is any finite or infinite subset of Act^* . For example, if $W = \{w \in \text{Act}^* \mid \text{action } \textit{crash} \text{ occurs in } w\}$, Safe_W corresponds to the property that action *crash* is never executed. We have to make sure in the first place that the sets Safe_W are measurable. We need the following lemma.

Lemma 5. *For each $W \subseteq \text{Act}^*$, Safe_W is measurable in Θ . Moreover, if W is finite, Safe_W can be written as a disjoint union of basic cylindrical sets.*

Corollary 2. *Let s be a state of L and $W \subseteq \text{Act}^*$. It holds that $\nu_s(\text{Safe}_W) = 1 - \mu_s(s \stackrel{W}{\Rightarrow})$.*

As a corollary of the previous result, we get the following, which characterizes \sqsubseteq in terms of probability of satisfying safety properties.

Theorem 3. *Let s, s' be states of L and suppose L has no dead states. We have $s \sqsubseteq s'$ if and only if for each $W \subseteq \text{Act}^*$, $\nu_s(\text{Safe}_W) \geq \nu_{s'}(\text{Safe}_W)$.*

Of course, \sqsubseteq can be also characterized in terms of *reachability* properties of the form $\{\tau \mid \tau \cap W \neq \emptyset\}$. In this case, the inequality between s and s' gets reversed.

6 Testing Separated RPLTS

In a separated system, actions can be partitioned into two sets: a set of actions Σ that are probabilistically controlled by the system, and a set of nondeterministic actions A that are under the control of the environment (observer). Accordingly, actions that are probabilistic for processes are nondeterministic for observers, and vice-versa. Importantly, the two types of actions do not mix up: the set of states as well gets partitioned into a set of states where only probabilistic choices are possible, and a set of states where only nondeterministic choices are possible. Nondeterministic actions can be modelled as actions that, if available, have probability one. These informal considerations lead to the next definition.




Fig. 2. Two separated RPLTS (left and center) and a non-separated one (right).


Definition 10 (separated processes and observers). Let (Σ, A) form a partition of Act , that is $\Sigma \cap A = \emptyset$ and $\Sigma \cup A = Act$. We say a RPLTS $L = (S, \delta, P)$ is a (Σ, A) -separated process if there is a partition of the states, $S = G \cup R$, such that

- for each $s \in G$ and $(s, a, s') \in \delta$ one has $a \in \Sigma$; moreover, if $(s, b, s'') \in \delta$, for some b and s'' , then $a = b$;
- for each $s \in R$ and $(s, a, s') \in \delta$ one has $a \in A$; moreover, if $(s, a, s'') \in \delta$, for some s'' , then $s' = s''$.

A (Σ, A) -separated observer is a $(A \cup \{\omega\}, \Sigma)$ -separated RPLTS, where ω is the distinct success action, $\omega \notin Act$.

Example 3. Let $A = \{b, c, d\}$ and $\Sigma = \{a, e\}$. An example of (Σ, A) -separated and one of non (Σ, A) -separated processes are depicted in Fig. 2.


Remark 2. Separated RPLTS are reminiscent of Hansson and Jonsson's *alternating model* [15], where probabilistic and nondeterministic states are strictly alternated with one another. In the alternating model, probabilistic choices are labeled by the silent action τ . In our model we do not have silent actions; but one can get rid of those τ 's by absorbing them into incoming edges of probabilistic states (see picture on the right). Modulo this transformation, separated RPLTS can be seen as a proper extension of the alternating model.



In what follows we fix a generic (Σ, A) -separated process L and a generic (Σ, A) -separated observer O . We let s be a state of L and o be a state of O . The proof of the main result in this section rests on the following crucial lemma.

Lemma 6. Consider $T \in \mathcal{M}_s^L$ and $U \in \mathcal{M}_o^O$. There is at most one $w \in Act^*$ such that $T \xrightarrow{w}$ and $U \xrightarrow{w\omega}$.

It is worthwhile to note that the above lemma fails to hold for non-separated system. As an example, consider T and U depicted on the right. Clearly, either T belongs to a $(\{c\}, \{a, b\})$ -separated RPLTS or U belongs to a $(\{a, b\}, \{c\})$ -separated observer (a and b are used in both as nondeterministic actions) and they violate Lemma 6.



Recall that $f_s^L : Act^* \rightarrow [0, 1]$ denotes the formal power series associated with the RPLTS L at state s . Similarly, $f_o^O : (Act \cup \{\omega\})^* \rightarrow [0, 1]$ is associated with the RPLTS O at state o .

Corollary 3. $\mu_{(s,o)}(\text{sat}(s,o)) = \sum_{w \in \text{Act}^*} f_s^L(w) \cdot f_o^O(w\omega)$.


As a consequence of Theorem 1 and Corollary 3 we get the following result.

Corollary 4 (coincidence of linear-time and separated may-testing semantics). *Let $\sqsubseteq^{(\Sigma,A)}$ be the may preorder on the states of L generated by (Σ, A) -separated observers. For each s and s' , one has $s \sqsubseteq^{(\Sigma,A)} s'$ if and only if $f_s^L \leq f_{s'}^L$.*

Example 4. Consider the RPLTS composed by the trees depicted in the left and in the center of Fig. 2. This is a $(\{a\}, \{b, c, d\})$ -separated RPLTS. It is easy to check that $s_0 \leq_{\text{lin}} q_0$. From Corollary 4 we get $s_0 \sqsubseteq^{(\{a\}, \{b, c, d\})} q_0$.

7 Conclusion and Related Works

There exist many formalisms for the specification of probabilistic processes and as many variants of probabilistic semantics. The conceptual difference between our approach and other testing theories has been discussed in the Introduction. Concerning the set-theoretical comparison between behavioural relations, we restrict our attention to one of the most recent proposals [7], and refer for the other approaches to the papers mentioned therein. Consider the pcsp processes $P = a.(b.c \frac{1}{2} \oplus b.d)$ and $Q = a.b.c \frac{1}{2} \oplus a.b.d$. The picture on the right is the description of P and Q 's operational semantics in terms RPLTS's (the operational model of [7] is in fact different from ours, because transitions lead to probability distributions on states, rather than states). P and Q are discriminated by the may preorder of [7], as witnessed by the test $s' = a.b.c.\omega \sqcap a.b.d.\omega$, which tells us that Q is not smaller than P . On the other hand, P and Q are equated by our may-preorder, which can be established by resorting to Theorem 1. This example shows that the presence of internal choice in [7] increases the distinguishing power of observers. Indeed, several of the axioms usually valid for classical csp (like, e.g., distributivity of prefixing w.r.t. internal choice) are no longer valid in pcsp. We conjecture that if the internal choice operator were taken out from csp, the may preorder of [7] would coincide with ours.



As one would expect, our may-testing equivalence is coarser than probabilistic bisimulation [18]. Indeed, any set W in the alternative characterization (or equivalently, any canonical observer O_W , see Theorem 1) can be characterized by a formula of the Probabilistic Modal Logic, which induces probabilistic bisimulation [18]. That the inclusion is strict is witnessed by processes s and s' above, which can be distinguished by probabilistic bisimulation.

As for future work, an obvious next-step is the study of *must*-testing behavioural relations. Also, decidability and algorithmic issues for the considered semantics deserve further investigation. Currently, the only known facts concern the linear-time setting: in the context of Rabin's probabilistic finite-state automata, which are equivalent to RPLTS, it is known that the preorder is undecidable, while the induced equivalence is decidable in polynomial time (see [20] and references therein).

References

1. Acciai, L., Boreale, M., De Nicola, R.: Linear-Time and May-Testing in a Probabilistic Reactive Setting. Full version available at <http://rap.dsi.unifi.it/~acciai/papers/prob-may.pdf>
2. Baier, C.: On the algorithmic verification of probabilistic systems. Universität Mannheim, Habilitation Thesis (1998)
3. Boreale, M., De Nicola, R., Pugliese, R.: Trace and Testing Equivalence in Asynchronous Processes. *Information and Computation* 172:139-164 (2002)
4. Chatzikokolakis, K., Palamidessi, C.: Making random choices invisible to the scheduler. In *Proc. of CONCUR*, LNCS 4703:42-58 (2007)
5. Cleaveland, R., Iyer, S.P., Narasimha, M.: Probabilistic temporal logics via the modal mu-calculus. *Theor. Comput. Sci.*, Volume 342(2-3) (2005)
6. Deng, D., van Glabbeek, R., Hennessy, M., Morgan, C.: Testing Finitary Probabilistic Processes. In *Proc. of CONCUR*, LNCS 5710:274-288 (2009)
7. Deng, D., van Glabbeek, R., Hennessy, M., Morgan, C.: Characterising testing preorders for finite probabilistic processes. *Logical Methods in Computer Science* 4(4):1-33 (2008)
8. De Nicola, R.: Extensional equivalences for transition systems. *Acta Informatica* 24(2):211-237 (1987)
9. De Nicola, R., Hennessy, M.: Testing equivalences for processes. *Theoretical Computer Science* 34:83-133 (1984)
10. Ésik, Z., Kuich, W.: Formal Tree Series. *Journal of Automata, Languages and Combinatorics* 8(2):219-285 (2003)
11. Georgievska, S., Andova, S.: Retaining the Probabilities in Probabilistic Testing Theory. In *Proc. of FOSSACS*, LNCS 6014:79-93 (2010)
12. van Glabbeek, R., Smolka, S., Steffen, B., Tofts, C.: Reactive, generative, and stratified models of probabilistic processes. *Information and Computation* 121(1):59-80 (1995)
13. van Glabbeek, R. J.: The linear time-branching time spectrum. in *Proc. of CONCUR*, LNCS 458:278-297 (1990)
14. Halmos, P.: *Measure theory*. Litton Educational Publishing, Inc. (1950)
15. Hansson, H., Jonsson, B.: A Calculus for Communicating Systems with Time and Probabilities. In *Proc. of IEEE Real-Time Systems Symposium*, pp. 278–287 (1990)
16. Hennessy, M., Milner, R.: Algebraic laws for nondeterminism and concurrency. *Journal of the ACM* 32(1):137-161 (1985)
17. Jonsson, B., Yi, W.: Testing Preorders for Probabilistic Processes can be Characterized by Simulations. *TCS* 282(1):33-51 (2002)
18. Larsen, K.G., Skou, A.: Bisimulation through Probabilistic Testing. *Inf. and Comp.* 94(1):1-28 (1991)
19. Segala, R.: Testing Probabilistic Automata. In *Proc. of CONCUR*, LNCS 1119:299-314 (1996)
20. Tzeng, W.-G.: A polynomial time algorithm for the equivalence of probabilistic automata. *SIAM Journal on Computing* 21(2):216-227 (1992)
21. Wang, Y., Larsen, K.G.: Testing Probabilistic and Nondeterministic Processes. In *Proc. of PSTV*, IFIP Transactions C-8:47-61 (1992)