Logical Methods in Computer Science
Vol. 10(1:16)2014, pp. 1–42
www.lmcs-online.org

Submitted Mar. 4, 2013
Published Mar. 3, 2014

REVISITING TRACE AND TESTING EQUIVALENCES FOR NONDETERMINISTIC AND PROBABILISTIC PROCESSES

MARCO BERNARDO [a], ROCCO DE NICOLA [b], AND MICHELE LORETI [c]

[a] Dipartimento di Scienze di Base e Fondamenti – Università di Urbino – Italy
e-mail address: marco.bernardo@uniurb.it

[b] IMT – Institute for Advanced Studies Lucca – Italy
e-mail address: rocco.denicola@imtlucca.it

[c] Dipartimento di Statistica, Informatica, Applicazioni – Università di Firenze – Italy
e-mail address: michele.loreti@unifi.it

Abstract. Two of the most studied extensions of trace and testing equivalences to nondeterministic and probabilistic processes induce distinctions that have been questioned and lack properties that are desirable. Probabilistic trace-distribution equivalence differentiates systems that can perform the same set of traces with the same probabilities, and is not a congruence for parallel composition. Probabilistic testing equivalence, which relies only on extremal success probabilities, is backward compatible with testing equivalences for restricted classes of processes, such as fully nondeterministic processes or generative/reactive probabilistic processes, only if specific sets of tests are admitted. In this paper, new versions of probabilistic trace and testing equivalences are presented for the general class of nondeterministic and probabilistic processes. The new trace equivalence is coarser because it compares execution probabilities of single traces instead of entire trace distributions, and turns out to be compositional. The new testing equivalence requires matching all resolutions of nondeterminism on the basis of their success probabilities, rather than comparing only extremal success probabilities, and considers success probabilities in a trace-by-trace fashion, rather than cumulatively on entire resolutions. It is fully backward compatible with testing equivalences for restricted classes of processes; as a consequence, the trace-by-trace approach uniformly captures the standard probabilistic testing equivalences for generative and reactive probabilistic processes. The paper discusses in full details the new equivalences and provides a simple spectrum that relates them with existing ones in the setting of nondeterministic and probabilistic processes.

1. Introduction

Modeling and abstraction are two key concepts of computer science that go hand in hand. If we wish to model a computer system for the purpose of (computer-aided) analysis, it is essential that the right level of abstraction is chosen when describing system behaviors. Operational models based on variants of automata or labeled transition systems very often

2012 ACM CCS: [**Theory of computation**]: Models of Computation—Probabilistic Computation; Models of computation—Concurrency—Process calculi.

Key words and phrases: Labeled Transition Systems, Probabilistic Models, Behavioral Equivalences.

LOGICAL METHODS IN COMPUTER SCIENCE
DOI:10.2168/LMCS-10(1:16)2014

provide descriptions that are too detailed; it is then necessary to resort to additional machineries to abstract from unwanted details. Behavioral equivalences are one of such machineries and indeed many equivalences have been proposed depending on the specific aspects of systems descriptions to ignore or the specific properties to capture. Equivalences are used to assess the relationships between different views of the same system. If both the specification and the implementation of a concurrent system are described via the same formalism, then the correctness of the latter with respect to the former can be established by studying their behavioral relationships.

Behavioral equivalences were first of all defined for labeled transition systems (LTS – set of states related via transitions each labeled with the action that gives rise to the state change [24]) that were used as models of nonderministic processes. Then, they have been extended/adapted to generalizations of such models to take into account probabilistic, stochastic, or timed behaviors.

Among the most important equivalences defined for abstracting unnecessary details of nondeterministic processes modeled as LTS, we would like to mention the following three:

- *trace equivalence*, equating systems performing the same sequences of actions,
- *testing equivalence*, equating systems reacting similarly to external experiments by peer systems, and
- *bisimulation equivalence*, equating systems performing the same sequences of actions and recursively exhibiting the same behavior after them.

Studies about their relationships have shown that the first equivalence is coarser than the second one, which in turn is coarser than the third one. A coarser equivalence provides a more abstract view of a system and produces more identifications.

When probabilities enter the game and probabilistic extensions of LTS are considered, the possible alternatives in choosing what to observe and compare, in deciding how to resolve nondeterminism, or in assembling the results of the observations are very many and the different choices can give rise to significantly different behavioral relations. Indeed, many proposals have been put forward and discussion is still going on about whether the identifications that these relations induce do capture the intuition one has in mind about the wanted behavior of probabilistic descriptions.

In this paper, we would like to concentrate on probabilistic trace and testing equivalences for processes described by means of an extension of the LTS model that combines nondeterminism and probabilities. The extended model, which we have thus called NPLTS, is such that every action-labeled transition goes from a source state to a probability distribution over target states – in the style of [26, 31] – rather than to a single target state.

The most used definition of probabilistic trace equivalence for nondeterministic and probabilistic processes is the one provided in [32]. To resolve nondeterminism, it resorts to the notion of *scheduler* (or adversary), which can be viewed as an external entity that selects the next action to perform according to the current state and the past history. When a scheduler is applied to a process, a fully probabilistic model called a *resolution* is obtained. Two processes are considered trace equivalent if, for each resolution of any of the two processes, there exists a resolution of the other process such that the probability of *each trace* is the same in the two resolutions. In other words, the two resolutions must exhibit the *same trace distribution*. We shall denote this equivalence by $\sim_{\mathrm{PTr,dis}}$.

Testing equivalence for the same class of processes has been studied in [39, 21, 33, 12]. It considers the probability of performing computations along which the same tests are passed, called successful computations. Due to the possible presence of equally labeled transitions departing from the same state, there is not necessarily a single probability value with which a nondeterministic and probabilistic process passes a test. Given two states s_1 and s_2 and the initial state o of an observer, this testing equivalence computes the probability of performing a successful computation from (s_1, o) and (s_2, o) in every maximal resolution of the interaction system resulting from the parallel composition of each process with the observer. Then, it compares *only extremal success probabilities*, i.e., the suprema ($\sqcup$) and the infima ($\sqcap$) of the success probabilities over all maximal resolutions of the two interaction systems. We shall denote this equivalence by $\sim_{\text{PTe-}\sqcup\sqcap}$.

After examining the above mentioned trace and testing equivalences for nondeterministic and probabilistic processes, we noticed that both equivalences induce differentiations that might be questionable and lack properties that are in general desirable.

For the equivalence $\sim_{\text{PTr,dis}}$, we have that it considers as inequivalent the two processes in Fig. 4 (p. 8), in spite of the fact that they can undoubtedly exhibit the same set of traces with the same probabilities. Moreover, $\sim_{\text{PTr,dis}}$ is not preserved by parallel composition. As shown in [32], given two $\sim_{\text{PTr,dis}}$-equivalent processes and given a third process, it is not necessarily the case that the parallel composition of the first process with the third one is $\sim_{\text{PTr,dis}}$-equivalent to the parallel composition of the second process with the third one.

The equivalence $\sim_{\text{PTe-}\sqcup\sqcap}$, instead, identifies the two processes in Fig. 5 (p. 16) mainly because its definition only considers extremal success probabilities. A consequence of such a choice is that this testing equivalence, contrary to what happens for the purely nondeterministic case, does not imply the trace equivalence $\sim_{\text{PTr,dis}}$. Indeed, the two processes in Fig. 5, which are identified by $\sim_{\text{PTe-}\sqcup\sqcap}$, are distinguished by $\sim_{\text{PTr,dis}}$. Actually, the inclusion depends on the type of schedulers used for deriving resolutions of nondeterminism; it holds if randomized schedulers are admitted for $\sim_{\text{PTr,dis}}$ as in [32], while it does not hold if only deterministic schedulers are considered.

Another characteristic of $\sim_{\text{PTe-}\sqcup\sqcap}$ is that of being only *partially* backward compatible with existing testing equivalences for restricted classes of processes. Compatibility depends on the set of admitted tests. For example, the two fully nondeterministic processes in Fig. 8 (p. 19) are identified by the original testing equivalence of [11]. The relation $\sim_{\text{PTe-}\sqcup\sqcap}$ equates them if only fully nondeterministic tests are employed, but distinguishes them as soon as probabilities are admitted within tests. Dually, following the terminology of [37], the two generative/reactive probabilistic processes in Fig. 10 (p. 20), which are identified by the generative probabilistic testing equivalence of [9] and the reactive probabilistic testing equivalence of [25], are equated by $\sim_{\text{PTe-}\sqcup\sqcap}$ if only generative/reactive probabilistic tests are employed, but are told apart by the same relation as soon as internal nondeterminism is admitted within tests.

Indeed, these two examples show that $\sim_{\text{PTe-}\sqcup\sqcap}$ is sensitive to the moment of occurrence of internal choices when testing fully nondeterministic processes (resp. generative/reactive probabilistic processes) with tests admitting probabilities (resp. internal nondeterminism), because it becomes possible to make copies of intermediate states of the processes under test. As pointed out in [1], this capability increases the distinguishing power of testing equivalence. In a probabilistic setting, this may lead to questionable estimations of success probabilities (see [16] and the references therein).

In this paper, we study new trace and testing equivalences (for nondeterministic and probabilistic processes) that, different from the old ones, do possess the above-mentioned properties. We shall start by defining a coarser probabilistic trace equivalence $\sim_{\text{PTr}}$ that, instead of considering entire trace distributions as in $\sim_{\mathrm{PTr,dis}}$, compares the execution probabilities of single traces. Moreover, we shall define a finer probabilistic testing equivalence $\sim_{\mathrm{PTe}\text{-}\forall\exists}$ that, instead of focussing only on the highest and the lowest probability of passing a test as in $\sim_{\mathrm{PTe}\text{-}\sqcup\sqcap}$, requires for each maximal resolution of the interaction system on one side the existence of a maximal resolution of the interaction system on the other side that has the same success probability.

While the new trace equivalence $\sim_{\mathrm{PTr}}$ reaches the goal of being compositional, the new testing equivalence $\sim_{\mathrm{PTe}\text{-}\forall\exists}$ is still not fully backward compatible with the testing equivalences for restricted classes of processes. We shall however use $\sim_{\mathrm{PTe}\text{-}\forall\exists}$ as a stepping stone to define another probabilistic testing equivalence, $\sim_{\mathrm{PTe\text{-}tbt}}$, that requires matching success probabilities of maximal resolutions in a *trace-by-trace* fashion rather than cumulatively over all successful computations of the maximal resolutions. This further testing equivalence is a fully conservative extension of the ones in [11, 9, 25] and avoids questionable estimations of success probabilities without resorting to model transformations as in [16]. Thus, the trace-by-trace approach provides a uniform way of defining testing equivalence over different probabilistic models. This means that the standard notions of testing equivalence for generative/reactive probabilistic processes could be redefined by following the same trace-by-trace approach taken for the general model, without altering their discriminating power. Interestingly, we shall see that $\sim_{\mathrm{PTe\text{-}tbt}}$ is comprised between $\sim_{\mathrm{PTr}}$ and a novel probabilistic failure equivalence $\sim_{\mathrm{PF}}$, which in turn is comprised between $\sim_{\mathrm{PTe\text{-}tbt}}$ and $\sim_{\mathrm{PTe}\text{-}\forall\exists}$.

For each of the equivalences considered in the paper, we shall introduce the two variants determined by the assumed nature of the schedulers used to resolve nondeterminism, namely *deterministic schedulers* or *randomized schedulers*.

The rest of the paper, which is a revised and extended version of [2], is organized as follows. Section 2 presents the necessary definitions for the NPLTS model. Section 3 introduces $\sim_{\mathrm{PTr}}$ and shows that it is a congruence with respect to parallel composition. Sections 4 and 5 respectively deal with $\sim_{\mathrm{PTe}\text{-}\forall\exists}$ and $\sim_{\mathrm{PTe\text{-}tbt}}$ by providing the necessary results to relate them to the new trace equivalence (inclusion) and to testing equivalences for restricted classes of processes (backward compatibility), emphasizing that the trace-by-trace approach unifies the testing equivalences defined for subclasses of NPLTS models without internal nondeterminism. Section 6 places in a spectrum old and new trace ad testing equivalences. Section 7 draws some conclusions and suggests future works.

2. Nondeterministic and Probabilistic Processes

Processes combining nondeterminism and probability are typically described by means of extensions of the LTS model, in which every action-labeled transition goes from a source state to a *probability distribution over target states* rather than to a single target state. They are essentially Markov decision processes [15] and are representative of a number of slightly different probabilistic computational models including internal nondeterminism that have appeared in the literature with names such as, e.g., concurrent Markov chains [38], alternating probabilistic models [18, 39, 29], NP-systems [20], probabilistic automata in the sense of [31], probabilistic processes in the sense of [21], denotational probabilistic models in the sense of [19], probabilistic transition systems in the sense of [22], and pLTS [12] (see [36] for an overview). We formalize them as a variant of simple probabilistic automata [31] and give them the acronym NPLTS to stress the possible simultaneous presence of nondeterminism (N) and probability (P) in the LTS-like model.

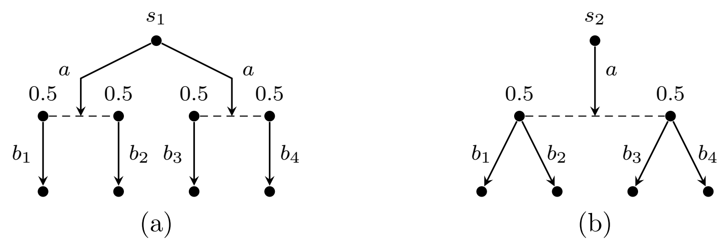

Figure 1: Graphical representation of two example NPLTS models

Definition 2.1. A *nondeterministic and probabilistic labeled transition system*, NPLTS for short, is a triple $(S, A, \longrightarrow)$ where:

- S is an at most countable set of states.
- A is a countable set of transition-labeling actions.
- $\longrightarrow \subseteq S \times A \times \mathit{Distr}(S)$ is a transition relation, where $\mathit{Distr}(S)$ is the set of discrete probability distributions over S. ■

A transition $(s, a, \mathcal{D})$ is written $s \xrightarrow{a} \mathcal{D}$. We say that $s' \in S$ is not reachable from s via that a-transition if $\mathcal{D}(s') = 0$, otherwise we say that it is reachable with probability $p = \mathcal{D}(s')$. The reachable states form the support of $\mathcal{D}$, i.e., $\mathit{supp}(\mathcal{D}) = \{s' \in S \mid \mathcal{D}(s') > 0\}$. The choice among all the transitions departing from s is nondeterministic and can be influenced by the external environment, while the choice of the target state for a specific transition is probabilistic and takes place internally.

An NPLTS can be depicted as a directed graph-like structure in which vertices represent states and action-labeled edges represent action-labeled transitions. Given a transition $s \xrightarrow{a} \mathcal{D}$, the corresponding a-labeled edge goes from the vertex for state s to a set of vertices linked by a dashed line, each of which represents a state $s' \in \mathit{supp}(\mathcal{D})$ and is labeled with $\mathcal{D}(s')$ – label omitted if $\mathcal{D}(s') = 1$. The graphical representation is exemplified in Fig. 1.

The NPLTS model embeds various less expressive models. In particular, it represents:

(1) A *fully nondeterministic process* when every transition leads to a distribution that concentrates all the probability mass into a single target state.
(2) A *fully probabilistic process* when every state has at most one outgoing transition.
(3) A *reactive probabilistic process* [37] – or probabilistic automaton in the sense of [30] – when no state has two or more outgoing transitions labeled with the same action.

The NPLTS in Fig. 1(a) mixes probability and internal nondeterminism, while the one in Fig. 1(b) describes a reactive probabilistic process. An example of fully probabilistic process can be obtained from the NPLTS in Fig. 1(a) by removing one of its two a-transitions.

In this setting, a computation is a sequence of state-to-state steps, each denoted by $s \overset{a}{\longmapsto} s'$ and derived from a state-to-distribution transition $s \xrightarrow{a} \mathcal{D}$.

Definition 2.2. Let $\mathcal{L} = (S, A, \longrightarrow)$ be an NPLTS and $s, s' \in S$. We say that:

$$c \equiv s_0 \overset{a_1}{\longmapsto} s_1 \overset{a_2}{\longmapsto} s_2 \dots s_{n-1} \overset{a_n}{\longmapsto} s_n$$

is a *computation* of $\mathcal{L}$ of length n from $s = s_0$ to $s' = s_n$ iff for all $i = 1, \dots, n$ there exists a transition $s_{i-1} \xrightarrow{a_i} \mathcal{D}_i$ such that $s_i \in \mathit{supp}(\mathcal{D}_i)$, with $\mathcal{D}_i(s_i)$ being the execution probability of step $s_{i-1} \overset{a_i}{\longmapsto} s_i$ conditioned on the selection of transition $s_{i-1} \xrightarrow{a_i} \mathcal{D}_i$ of $\mathcal{L}$ at state s_{i-1}. We say that c is *maximal* iff it is not a proper prefix of any other computation from s. We denote by $\mathcal{C}_{\text{fin}}(s)$ the set of finite-length computations from s. ■

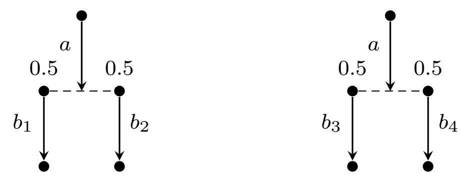

Figure 2: The two maximal resolutions of the NPLTS in Fig. 1(a)

A resolution of a state s of an NPLTS $\mathcal{L}$ is the result of any possible way of resolving nondeterminism starting from s. A resolution is a tree-like structure whose branching points represent probabilistic choices. This is obtained by unfolding from s the graph structure underlying $\mathcal{L}$ and by selecting at each state a single transition of $\mathcal{L}$ (*deterministic scheduler*) or a combined transition of $\mathcal{L}$ (*randomized scheduler*) among all the transitions that are possible from the reached state. We shall consider only history-independent schedulers.

Below, we formalize the notion of resolution arising from a deterministic scheduler as a fully probabilistic NPLTS. Notice that, when $\mathcal{L}$ is fully nondeterministic, these resolutions coincide with the computations of $\mathcal{L}$.

Definition 2.3. Let $\mathcal{L} = (S, A, \longrightarrow)$ be an NPLTS and $s \in S$. We say that an NPLTS $\mathcal{Z} = (Z, A, \longrightarrow_{\mathcal{Z}})$ is a *resolution* of s obtained via a *deterministic scheduler* iff there exists a state correspondence function $corr_{\mathcal{Z}} : Z \to S$ such that $s = corr_{\mathcal{Z}}(z_s)$, for some $z_s \in Z$, and for all $z \in Z$ it holds that:

- If $z \xrightarrow{a}_{\mathcal{Z}} \mathcal{D}$, then $corr_{\mathcal{Z}}(z) \xrightarrow{a} \mathcal{D}'$ with $\mathcal{D}(z') = \mathcal{D}'(corr_{\mathcal{Z}}(z'))$ for all $z' \in Z$.
- If $z \xrightarrow{a_1}_{\mathcal{Z}} \mathcal{D}_1$ and $z \xrightarrow{a_2}_{\mathcal{Z}} \mathcal{D}_2$, then $a_1 = a_2$ and $\mathcal{D}_1 = \mathcal{D}_2$.

We say that $\mathcal{Z}$ is *maximal* iff it cannot be further extended in accordance with the graph structure of $\mathcal{L}$ and the constraints above. We denote by $Res(s)$ and $Res_{\max}(s)$ the sets of resolutions and maximal resolutions of s obtained via deterministic schedulers. ■

Since $\mathcal{Z} \in Res(s)$ is fully probabilistic, the probability $prob(c)$ of executing $c \in \mathcal{C}_{\text{fin}}(z_s)$ can be defined as the product of the (no longer conditional) execution probabilities of the individual steps of c, with $prob(c)$ being always equal to 1 if $\mathcal{L}$ is fully nondeterministic. This notion is lifted to $\mathcal{C} \subseteq \mathcal{C}_{\text{fin}}(z_s)$ by letting $prob(\mathcal{C}) = \sum_{c \in \mathcal{C}} prob(c)$ whenever none of the computations in $\mathcal{C}$ is a proper prefix of one of the others. The two maximal resolutions of the NPLTS in Fig. 1(a) are shown in Fig. 2; both of them possess two maximal computations, each having probability 0.5.

The transitions of a resolution obtained via a randomized scheduler are not necessarily ordinary transitions of $\mathcal{L}$, but combined transitions derived as convex combinations of equally labeled transitions of the original model. Formally, the first clause of Def. 2.3 changes as follows:

- If $z \xrightarrow{a}_{\mathcal{Z}} \mathcal{D}$, then there are $n \in \mathbb{N}_{>0}$, $(p_i \in \mathbb{R}_{]0,1]} \mid 1 \leq i \leq n)$, and $(corr_{\mathcal{Z}}(z) \xrightarrow{a} \mathcal{D}_i \mid 1 \leq i \leq n)$ such that $\sum_{i=1}^{n} p_i = 1$ and $\mathcal{D}(z') = \sum_{i=1}^{n} p_i \cdot \mathcal{D}_i(corr_{\mathcal{Z}}(z'))$ for all $z' \in Z$.

It is worth noting that an ordinary transition is a combined transition in which $n = 1$ and $p_1 = 1$ and that, when $\mathcal{L}$ has no internal nondeterminism (like in the fully/reactive probabilistic case), a resolution arising from randomized schedulers can only be originated by a convex combination of a transition with itself. In the following, we use the shorthand ct for "based on combined transitions". We thus denote by $Res^{\text{ct}}(s)$ and $Res^{\text{ct}}_{\max}(s)$ the sets of resolutions and maximal resolutions of s obtained via randomized schedulers.

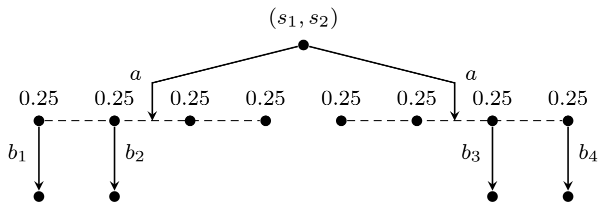

Figure 3: Fully synchronous parallel composition of the two NPLTS models in Fig. 1

We finally introduce a parallel operator $\_ \parallel_{\mathcal{A}} \_$ for NPLTS models that synchronize on a set of actions $\mathcal{A}$ and proceed independently of each other on any other action. The adoption of a CSP-like parallel operator is by now standard in the definition of testing equivalences for probabilistic processes (see, e.g., [21, 33, 9, 12]). We have preferred using this operator rather than a CCS-like parallel operator because the former embodies a mechanism for enforcing synchronizations, while the latter does not and hence, when defining testing equivalences, requires either resorting to an additional operator (e.g., restriction in a CCS setting as in [39]) or considering only computations whose steps are all labeled with invisible τ-actions stemming from the synchronization of an action with the corresponding coaction (like in traditional testing theory [11]). We would, however, like to stress that, if we had used a CCS-like parallel composition supporting τ-labeled two-way synchronizations, the resulting testing equivalences and the compositionality results would have been much the same.

Definition 2.4. Let $\mathcal{L}_i = (S_i, A, \longrightarrow_i)$ be an NPLTS for $i = 1, 2$ and $\mathcal{A} \subseteq A$. The *parallel composition* of $\mathcal{L}_1$ and $\mathcal{L}_2$ with synchronization on $\mathcal{A}$ is the NPLTS $\mathcal{L}_1 \parallel_{\mathcal{A}} \mathcal{L}_2 = (S_1 \times S_2, A, \longrightarrow)$ where $\longrightarrow \subseteq (S_1 \times S_2) \times A \times \mathit{Distr}(S_1 \times S_2)$ is such that $(s_1, s_2) \xrightarrow{a} \mathcal{D}$ iff one of the following holds:

- $a \in \mathcal{A}$, $s_1 \xrightarrow{a}_1 \mathcal{D}_1$, $s_2 \xrightarrow{a}_2 \mathcal{D}_2$, and $\mathcal{D}(s'_1, s'_2) = \mathcal{D}_1(s'_1) \cdot \mathcal{D}_2(s'_2)$ for all $(s'_1, s'_2) \in S_1 \times S_2$.
- $a \notin \mathcal{A}$, $s_1 \xrightarrow{a}_1 \mathcal{D}_1$, $\mathcal{D}(s'_1, s'_2) = \mathcal{D}_1(s'_1)$ if $s'_2 = s_2$, and $\mathcal{D}(s'_1, s'_2) = 0$ if $s'_2 \in S_2 \setminus \{s_2\}$.
- $a \notin \mathcal{A}$, $s_2 \xrightarrow{a}_2 \mathcal{D}_2$, $\mathcal{D}(s'_1, s'_2) = \mathcal{D}_2(s'_2)$ if $s'_1 = s_1$, and $\mathcal{D}(s'_1, s'_2) = 0$ if $s'_1 \in S_1 \setminus \{s_1\}$. ∎

Throughout the paper, we shall use $\mathcal{L}_1 \parallel \mathcal{L}_2$ to denote the fully synchronous parallel composition $\mathcal{L}_1 \parallel_A \mathcal{L}_2$. Figure 3 shows the NPLTS resulting from the fully synchronous parallel composition of the two NPLTS models in Fig. 1. Note that the two nondeterministic choices after the a-transition of the NPLTS in Fig. 1(b) have disappeared in Fig. 3, because the synchronization between a state with a single transition and a state with several differently labeled transitions always results in a state with at most a single transition.

3. Trace Equivalences for NPLTS Models

Trace equivalences for NPLTS models examine the probability with which two states perform computations labeled with the same action sequences, called traces, for each possible way of resolving nondeterminism. We say that a finite-length computation is *compatible* with a trace $\alpha \in A^*$ iff the sequence of actions labeling the computation steps is equal to α. Given an NPLTS $\mathcal{L} = (S, A, \longrightarrow)$ and a resolution $\mathcal{Z}$ of a state s, we denote by $\mathcal{CC}(z_s, \alpha)$ the set of α-compatible computations in $\mathcal{C}_{\mathrm{fin}}(z_s)$. We now recall two variants of the probabilistic trace-distribution equivalence introduced in [32] and further studied in [7, 27, 28, 6].

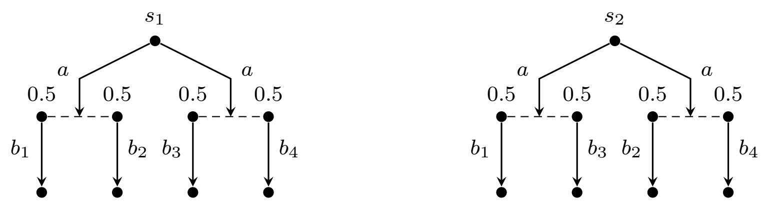

Figure 4: NPLTS models distinguished by $\sim_{\text{PTr,dis}}/\sim_{\text{PTr,dis}}^{\text{ct}}$ and identified by $\sim_{\text{PTr}}/\sim_{\text{PTr}}^{\text{ct}}$

Definition 3.1. Let $(S, A, \longrightarrow)$ be an NPLTS. We say that $s_1, s_2 \in S$ are *probabilistic trace-distribution equivalent*, written $s_1 \sim_{\text{PTr,dis}} s_2$, iff:

- For each $\mathcal{Z}_1 \in \mathit{Res}(s_1)$ there exists $\mathcal{Z}_2 \in \mathit{Res}(s_2)$ such that $\underline{\text{for all } \alpha \in A^*}$:
$$\mathit{prob}(\mathcal{CC}(z_{s_1}, \alpha)) \;=\; \mathit{prob}(\mathcal{CC}(z_{s_2}, \alpha))$$
- For each $\mathcal{Z}_2 \in \mathit{Res}(s_2)$ there exists $\mathcal{Z}_1 \in \mathit{Res}(s_1)$ such that $\underline{\text{for all } \alpha \in A^*}$:
$$\mathit{prob}(\mathcal{CC}(z_{s_2}, \alpha)) \;=\; \mathit{prob}(\mathcal{CC}(z_{s_1}, \alpha))$$

We denote by $\sim_{\text{PTr,dis}}^{\text{ct}}$ the coarser variant based on randomized schedulers. ■

The relations $\sim_{\text{PTr,dis}}$ and $\sim_{\text{PTr,dis}}^{\text{ct}}$ are quite discriminating because they compare entire trace distributions and hence impose a constraint on the execution probability of *all the traces* of any pair of corresponding resolutions (*fully matching resolutions*). For instance, states s_1 and s_2 in Fig. 4 are distinguished by $\sim_{\text{PTr,dis}}$ because neither of the two maximal resolutions of s_1, which are depicted in Fig. 2, is matched according to Def. 3.1 by (i.e., has the same trace distribution as) one of the two maximal resolutions of s_2.

However, s_1 and s_2 have exactly the same set of traces, which is $\{\varepsilon, a, a\,b_1, a\,b_2, a\,b_3, a\,b_4\}$, and each of these traces has the same probability of being performed in both processes once nondeterminism has been resolved, hence it might seem reasonable to identify s_1 and s_2. The constraint on trace distributions can indeed be relaxed by considering *a single trace at a time*, i.e., by anticipating the quantification over traces with respect to the quantification over resolutions in Def. 3.1. In this way, differently labeled computations of a resolution are allowed to be matched by computations of different resolutions (*partially matching resolutions*), which leads to the following new probabilistic trace equivalences.

Definition 3.2. Let $(S, A, \longrightarrow)$ be an NPLTS. We say that $s_1, s_2 \in S$ are *probabilistic trace equivalent*, written $s_1 \sim_{\text{PTr}} s_2$, iff $\underline{\text{for all } \alpha \in A^*}$ it holds that:

- For each $\mathcal{Z}_1 \in \mathit{Res}(s_1)$ there exists $\mathcal{Z}_2 \in \mathit{Res}(s_2)$ such that:
$$\mathit{prob}(\mathcal{CC}(z_{s_1}, \alpha)) \;=\; \mathit{prob}(\mathcal{CC}(z_{s_2}, \alpha))$$
- For each $\mathcal{Z}_2 \in \mathit{Res}(s_2)$ there exists $\mathcal{Z}_1 \in \mathit{Res}(s_1)$ such that:
$$\mathit{prob}(\mathcal{CC}(z_{s_2}, \alpha)) \;=\; \mathit{prob}(\mathcal{CC}(z_{s_1}, \alpha))$$

We denote by $\sim_{\text{PTr}}^{\text{ct}}$ the coarser variant based on randomized schedulers. ■

Theorem 3.3. *Let $(S, A, \longrightarrow)$ be an NPLTS and $s_1, s_2 \in S$. Then:*
$$s_1 \sim_{\text{PTr,dis}} s_2 \implies s_1 \sim_{\text{PTr}} s_2$$
$$s_1 \sim_{\text{PTr,dis}}^{\text{ct}} s_2 \implies s_1 \sim_{\text{PTr}}^{\text{ct}} s_2$$

Proof. If $s_1 \sim_{\text{PTr,dis}} s_2$ (resp. $s_1 \sim_{\text{PTr,dis}}^{\text{ct}} s_2$), then $s_1 \sim_{\text{PTr}} s_2$ (resp. $s_1 \sim_{\text{PTr}}^{\text{ct}} s_2$) follows by taking the same fully matching resolutions considered for $\sim_{\text{PTr,dis}}$ (resp. $\sim_{\text{PTr,dis}}^{\text{ct}}$). □

The implications in Thm. 3.3 cannot be reversed. For example, in Fig. 4 it holds that $s_1 \sim_{\rm PTr} s_2$ because the leftmost maximal resolution of s_1 is matched by the leftmost maximal resolution of s_2 with respect to trace $a\, b_1$, and by the rightmost maximal resolution of s_2 with respect to trace $a\, b_2$. Figures 4 and 5 (p. 16) together show that $\sim_{\rm PTr}$ and $\sim_{\rm PTr,dis}^{\rm ct}$ are incomparable with each other.

All the four trace equivalences above are fully backward compatible with the two trace equivalences respectively defined in [5] for fully nondeterministic processes – denoted by $\sim_{\rm Tr,fnd}$ – and in [23] for fully probabilistic processes – denoted by $\sim_{\rm Tr,fpr}$. Moreover, they are partially backward compatible with the trace equivalence – denoted by $\sim_{\rm Tr,rpr}$ – that can be defined for reactive probabilistic processes by following one of the approaches in [35].

Theorem 3.4. *Let $\mathcal{L} = (S, A, \longrightarrow)$ be an NPLTS and $s_1, s_2 \in S$.*

(1) *If $\mathcal{L}$ is fully nondeterministic, then:*

$$s_1 \sim_{\rm PTr,dis} s_2 \;\Leftrightarrow\; s_1 \sim_{\rm PTr,dis}^{\rm ct} s_2 \;\Leftrightarrow\; s_1 \sim_{\rm PTr} s_2 \;\Leftrightarrow\; s_1 \sim_{\rm PTr}^{\rm ct} s_2 \;\Leftrightarrow\; s_1 \sim_{\rm Tr,fnd} s_2$$

(2) *If $\mathcal{L}$ is fully probabilistic, then:*

$$s_1 \sim_{\rm PTr,dis} s_2 \;\Leftrightarrow\; s_1 \sim_{\rm PTr,dis}^{\rm ct} s_2 \;\Leftrightarrow\; s_1 \sim_{\rm PTr} s_2 \;\Leftrightarrow\; s_1 \sim_{\rm PTr}^{\rm ct} s_2 \;\Leftrightarrow\; s_1 \sim_{\rm Tr,fpr} s_2$$

(3) *If $\mathcal{L}$ is reactive probabilistic, then:*

$$\begin{aligned} s_1 \sim_{\rm PTr,dis} s_2 &\Rightarrow s_1 \sim_{\rm Tr,rpr} s_2 \\ s_1 \sim_{\rm PTr,dis}^{\rm ct} s_2 &\Rightarrow s_1 \sim_{\rm Tr,rpr} s_2 \\ s_1 \sim_{\rm PTr} s_2 &\Rightarrow s_1 \sim_{\rm Tr,rpr} s_2 \\ s_1 \sim_{\rm PTr}^{\rm ct} s_2 &\Rightarrow s_1 \sim_{\rm Tr,rpr} s_2 \end{aligned}$$

Proof. We proceed as follows:

(1) Suppose that $\mathcal{L}$ is fully nondeterministic. We recall from [5] that $s_1 \sim_{\rm Tr,fnd} s_2$ means that, for all $\alpha \in A^*$, there is an α-compatible computation from s_1 iff there is an α-compatible computation from s_2. The result is a straightforward consequence of the fact that the resolutions of $\mathcal{L}$ coincide with the computations of $\mathcal{L}$, hence the probability of performing within a resolution of $\mathcal{L}$ a computation compatible with a given trace can only be 1 or 0. Note that randomized schedulers are not important in this setting because, due to the absence of probabilistic choices, the model cannot contain submodels that arise from convex combinations of other submodels.

(2) Suppose that $\mathcal{L}$ is fully probabilistic. We recall from [23] that $s_1 \sim_{\rm Tr,fpr} s_2$ means that, for all $\alpha \in A^*$, $prob(\mathcal{CC}(s_1, \alpha)) = prob(\mathcal{CC}(s_2, \alpha))$. The result is a straightforward consequence of the fact that $\mathcal{L}$ has a single maximal resolution, which coincides with $\mathcal{L}$ itself. Note that schedulers are not important in this setting because there is no nondeterminism.

(3) Suppose that $\mathcal{L}$ is reactive probabilistic. Due to the absence of internal nondeterminism, $\sim_{\rm Tr,rpr}$ can be defined in the same way as $\sim_{\rm Tr,fpr}$ provided that, given $\alpha \in A^*$, probabilities of the form $prob(\mathcal{CC}(s, \alpha))$ are viewed as being conditional [35] on selecting the maximal resolution of $s \in S$ that contains all the α-compatible computations from s (this resolution is unique because $\mathcal{L}$ is reactive probabilistic). The result immediately follows. □

In the reactive probabilistic case, the first two implications cannot be reversed. If we consider a variant of s_1 (resp. s_2) in Fig. 4 having a single outgoing a-transition reaching with probability 0.5 a state with a b_1-transition and a b_2-transition (resp. b_3-transition) and with probability 0.5 a state with a b_3-transition (resp. b_2-transition) and a b_4-transition, then the two resulting states are related by $\sim_{\rm Tr,rpr}$ but distinguished by $\sim_{\rm PTr,dis}$ and $\sim_{\rm PTr,dis}^{\rm ct}$.

Interestingly, $\sim_{\rm PTr}$ and $\sim_{\rm PTr}^{\rm ct}$ are congruences with respect to parallel composition. This is quite surprising because, while $\sim_{\rm Tr,fnd}$ is compositional [5], all probabilistic trace semantics proposed so far in the literature, i.e., $\sim_{\rm Tr,fpr}$, $\sim_{\rm PTr,dis}$, and $\sim_{\rm PTr,dis}^{\rm ct}$, are *not* compositional [23, 32]. In particular, in [27] it was shown that the coarsest congruence contained in $\sim_{\rm PTr,dis}^{\rm ct}$ is a variant of the simulation equivalence of [34], while in [6] distributed schedulers (as opposed to centralized ones) were introduced to achieve compositionality.

To prove preservation of $\sim_{\rm PTr}$ under parallel composition, we make use of an alternative characterization of $\sim_{\rm PTr}$ itself based on *weighted traces*, each of which is an element of $A^* \times \mathbb{R}_{]0,1]}$. Before defining the function that associates the set of its weighted traces with each state, we introduce the following auxiliary notation where $X, Y \subseteq A^* \times \mathbb{R}_{]0,1]}$, $a \in A$, $\alpha \in A^*$, $p \in \mathbb{R}_{]0,1]}$, and $q \in \mathbb{R}_{[0,1]}$:

- $X \vdash (\alpha, q)$ iff either $(\alpha, q) \in X$, or $q = 0$ and $(\alpha, p') \notin X$ for all $p' \in \mathbb{R}_{]0,1]}$.
- $X + Y \ = \ \{(\alpha, q_1 + q_2) \mid X \vdash (\alpha, q_1) \wedge Y \vdash (\alpha, q_2) \wedge q_1 + q_2 > 0\}$.
- $a.X \ = \ \{(a\,\alpha, p') \mid (\alpha, p') \in X\}$.
- $p \cdot X \ = \ \{(\alpha, p \cdot p') \mid (\alpha, p') \in X\}$.

Definition 3.5. Let $(S, A, \longrightarrow)$ be an NPLTS. The set of functions $traces_i : S \rightarrow 2^{A^* \times \mathbb{R}_{]0,1]}}$, $i \in \mathbb{N}$, is inductively defined as follows:

- $traces_0(s) \ = \ \{(\varepsilon, 1)\}$.
- $traces_{i+1}(s) \ = \ \{(\varepsilon, 1)\} \cup \bigcup\limits_{s \xrightarrow{a} \mathcal{D}} a.(\sum\limits_{s' \in supp(\mathcal{D})} \mathcal{D}(s') \cdot traces_i(s'))$.

We let $traces(s) \ = \ \bigcup\limits_{i \in \mathbb{N}} traces_i(s)$. ■

For every $i \in \mathbb{N}$, function $traces_i$ maps each state s to the set of weighted traces built by considering only the computations from s of length at most i. The set $traces(s)$ is then obtained by considering all finite-length computations from s. The following lemma guarantees that the construction is monotonic.

Lemma 3.6. *Let $(S, A, \longrightarrow)$ be an NPLTS. For all $s \in S$ and $i \in \mathbb{N}$ it holds that:*

$$traces_i(s) \ \subseteq \ traces_{i+1}(s)$$

Proof. We prove that for all $s \in S$, $i \in \mathbb{N}$, $\alpha \in A^*$, and $p \in \mathbb{R}_{]0,1]}$ it holds that $(\alpha, p) \in traces_i(s)$ implies $(\alpha, p) \in traces_{i+1}(s)$ by proceeding by induction on the length of α.
Base of induction: Let $|\alpha| = 0$, i.e., $\alpha = \varepsilon$. Directly from Def. 3.5, for all $j \in \mathbb{N}$ we have that $(\varepsilon, p) \in traces_j(s)$ iff $p = 1$. Hence, the result holds when $\alpha = \varepsilon$.
Induction hypothesis: We assume that for all $s' \in S$, $j \in \mathbb{N}$, $\alpha' \in A^*$, and $p' \in \mathbb{R}_{]0,1]}$ it holds that $(\alpha', p') \in traces_j(s')$ implies $(\alpha', p') \in traces_{j+1}(s')$ when $|\alpha'| \leq n$ for some $n \in \mathbb{N}$.
Induction step: Let $\alpha = a\,\alpha'$ with $|\alpha'| = n$ and suppose that $(\alpha, p) \in traces_i(s)$. Then there exists a transition $s \xrightarrow{a} \mathcal{D}$ such that:

$$(\alpha', p) \ \in \ \sum_{s' \in supp(\mathcal{D})} \mathcal{D}(s') \cdot traces_{i-1}(s')$$

Hence, for each $s' \in supp(\mathcal{D})$ there exists $p_{s'} \in \mathbb{R}_{[0,1]}$ such that $traces_{i-1}(s') \vdash (\alpha', p_{s'})$, and $p = \sum_{s' \in supp(\mathcal{D})} \mathcal{D}(s') \cdot p_{s'}$. By the induction hypothesis, we have that $(\alpha', p_{s'}) \in traces_i(s')$ for each $s' \in supp(\mathcal{D})$ such that $(\alpha', p_{s'}) \in traces_{i-1}(s')$. Therefore:

$$(\alpha', p) \ \in \ \sum_{s' \in supp(\mathcal{D})} \mathcal{D}(s') \cdot traces_i(s')$$

which implies $(\alpha, p) \in traces_{i+1}(s)$. □

We now show that function *traces* can be used to provide an alternative definition of $\sim_{\mathrm{PTr}}$, which will be exploited at the end of this section to prove that $\sim_{\mathrm{PTr}}$ is preserved under parallel composition. The key property is that (α, p) is a weighted trace associated with a state s iff there exists a resolution of s where trace α can occur with probability p.

Lemma 3.7. *Let $(S, A, \longrightarrow)$ be an NPLTS. For all $s \in S$, $\alpha \in A^*$, and $p \in \mathbb{R}_{]0,1]}$ it holds that:*

$$(\alpha, p) \in \mathit{traces}(s) \iff \exists \mathcal{Z} \in \mathit{Res}(s).\, \mathit{prob}(\mathcal{CC}(z_s, \alpha)) = p$$

Proof. We prove the result by proceeding by induction on the length of α.
Base of induction: Let $|\alpha| = 0$, i.e., $\alpha = \varepsilon$. Directly from Def. 3.5, for all $j \in \mathbb{N}$ we have that $(\varepsilon, p) \in \mathit{traces}_j(s)$ iff $p = 1$. Moreover, for each $\mathcal{Z} \in \mathit{Res}(s)$ it holds that $\mathit{prob}(\mathcal{CC}(z_s, \varepsilon)) = 1$. Hence, the result holds when $\alpha = \varepsilon$.
Induction hypothesis: We assume that for all $s' \in S$, $\alpha' \in A^*$, and $p' \in \mathbb{R}_{]0,1]}$ it holds that $(\alpha', p') \in \mathit{traces}(s')$ iff there exists $\mathcal{Z} \in \mathit{Res}(s')$ such that $\mathit{prob}(\mathcal{CC}(z_{s'}, \alpha')) = p'$ when $|\alpha'| \leq n$ for some $n \in \mathbb{N}$.
Induction step: Let $\alpha = a\,\alpha'$ with $|\alpha'| = n$. Suppose that $(\alpha, p) \in \mathit{traces}(s)$. This means that $(\alpha, p) \in \mathit{traces}_i(s)$ for some $i \in \mathbb{N}$. Then there exists a transition $s \xrightarrow{a} \mathcal{D}$ such that:

$$(\alpha', p) \;\in \sum_{s' \in \mathit{supp}(\mathcal{D})} \mathcal{D}(s') \cdot \mathit{traces}_{i-1}(s')$$

Hence, for each $s' \in \mathit{supp}(\mathcal{D})$ there exists $p_{s'} \in \mathbb{R}_{[0,1]}$ such that $\mathit{traces}_{i-1}(s') \vdash (\alpha', p_{s'})$, and $p = \sum_{s' \in \mathit{supp}(\mathcal{D})} \mathcal{D}(s') \cdot p_{s'}$. Since $\mathit{traces}_{i-1}(s') \subseteq \mathit{traces}(s')$, by the induction hypothesis we have that there exists $\mathcal{Z}_{s'} \in \mathit{Res}(s')$ such that $\mathit{prob}(\mathcal{CC}(z_{s'}, \alpha')) = p_{s'}$ for each $s' \in \mathit{supp}(\mathcal{D})$ such that $(\alpha', p_{s'}) \in \mathit{traces}_{i-1}(s')$. Therefore, if we consider the resolution $\mathcal{Z} \in \mathit{Res}(s)$ that first selects transition $s \xrightarrow{a} \mathcal{D}$ and then behaves as $\mathcal{Z}_{s'}$ for each $s' \in \mathit{supp}(\mathcal{D})$ such that $(\alpha', p_{s'}) \in \mathit{traces}_{i-1}(s')$ whereas it halts in each $s' \in \mathit{supp}(\mathcal{D})$ such that $(\alpha', p_{s'}) \notin \mathit{traces}_{i-1}(s')$, it is easy to see that $\mathit{prob}(\mathcal{CC}(z_s, \alpha)) = p$.
Suppose now that there exists $\mathcal{Z} = (Z, A, \longrightarrow_{\mathcal{Z}}) \in \mathit{Res}(s)$ such that $\mathit{prob}(\mathcal{CC}(z_s, \alpha)) = p$. Then there exists a transition $z_s \xrightarrow{a}_{\mathcal{Z}} \mathcal{D}$ such that:

$$p \;= \sum_{z' \in \mathit{supp}(\mathcal{D})} \mathcal{D}(z') \cdot \mathit{prob}(\mathcal{CC}(z', \alpha'))$$

Hence, for each $z' \in \mathit{supp}(\mathcal{D})$ there exists $p_{z'} \in \mathbb{R}_{[0,1]}$ such that $p_{z'} = \mathit{prob}(\mathcal{CC}(z', \alpha'))$, and $p = \sum_{z' \in \mathit{supp}(\mathcal{D})} \mathcal{D}(z') \cdot p_{z'}$. Denoting by $\mathit{corr}_{\mathcal{Z}}$ the correspondence function for $\mathcal{Z}$, by the induction hypothesis we have that $(\alpha', p_{z'}) \in \mathit{traces}(\mathit{corr}_{\mathcal{Z}}(z'))$ for each $z' \in \mathit{supp}(\mathcal{D})$ such that $p_{z'} > 0$. Due to Lemma 3.6, for all $i \in \mathbb{N}_{\geq |\alpha'|}$ it holds that $(\alpha', p_{z'}) \in \mathit{traces}_i(\mathit{corr}_{\mathcal{Z}}(z'))$ for each $z' \in \mathit{supp}(\mathcal{D})$ such that $p_{z'} > 0$. Since there must exist a transition $s \xrightarrow{a} \mathcal{D}'$ such that $\mathcal{D}(z') = \mathcal{D}'(\mathit{corr}_{\mathcal{Z}}(z'))$ for all $z' \in Z$, it holds that:

$$(\alpha', p) \;\in \sum_{\mathit{corr}_{\mathcal{Z}}(z') \in \mathit{supp}(\mathcal{D}')} \mathcal{D}'(\mathit{corr}_{\mathcal{Z}}(z')) \cdot \mathit{traces}_{|\alpha'|}(\mathit{corr}_{\mathcal{Z}}(z'))$$

and hence:

$$(\alpha, p) \;\in\; a.(\sum_{\mathit{corr}_{\mathcal{Z}}(z') \in \mathit{supp}(\mathcal{D}')} \mathcal{D}'(\mathit{corr}_{\mathcal{Z}}(z')) \cdot \mathit{traces}_{|\alpha'|}(\mathit{corr}_{\mathcal{Z}}(z'))) \;\subseteq\; \mathit{traces}_{|\alpha|}(s)$$

which implies $(\alpha, p) \in \mathit{traces}(s)$. ∎

Theorem 3.8. *Let $(S, A, \longrightarrow)$ be an NPLTS and $s_1, s_2 \in S$. Then:*

$$s_1 \sim_{\mathrm{PTr}} s_2 \iff \mathit{traces}(s_1) = \mathit{traces}(s_2)$$

Proof. Directly from Def. 3.2 and Lemma 3.7. Notice that, given $\alpha \in A^*$, from the point of view of $\sim_{\mathrm{PTr}}$ a resolution $\mathcal{Z} \in \mathit{Res}(s_k)$, $k = 1, 2$, such that $\mathit{prob}(\mathcal{CC}(z_{s_k}, \alpha)) = 0$ is always matched by the resolution of s_{3-k} having only the initial state. Therefore, the exclusion of weighted traces with weight 0 from the set resulting from the application of function *traces* does not violate the present characterization of $\sim_{\mathrm{PTr}}$. $\square$

We finally exploit the result in Thm. 3.8 to show that $\sim_{\mathrm{PTr}}$ is preserved under parallel composition. This is an important and much wanted property that is essential for behavioral equivalences to support the compositional analysis of system descriptions.

Theorem 3.9. *Let $\mathcal{L}_k = (S_k, A, \longrightarrow_k)$ be an NPLTS for $k = 0, 1, 2$ and consider $\mathcal{L}_1 \parallel_{\mathcal{A}} \mathcal{L}_0$ and $\mathcal{L}_2 \parallel_{\mathcal{A}} \mathcal{L}_0$ for $\mathcal{A} \subseteq A$. Let $s_k \in S_k$ for $k = 0, 1, 2$. Then:*

$$s_1 \sim_{\mathrm{PTr}} s_2 \implies (s_1, s_0) \sim_{\mathrm{PTr}} (s_2, s_0)$$

Proof. For $\alpha_1, \alpha_2, \alpha \in A^*$, we let $\alpha_1 \otimes_{\mathcal{A}} \alpha_2 \vdash \alpha$ denote the smallest relation induced by the following inference rules:

$$\frac{}{\varepsilon \otimes_{\mathcal{A}} \varepsilon \vdash \varepsilon} \quad \frac{\alpha_1 \otimes_{\mathcal{A}} \alpha_2 \vdash \alpha}{a\,\alpha_1 \otimes_{\mathcal{A}} a\,\alpha_2 \vdash a\,\alpha}\ a \in \mathcal{A} \quad \frac{\alpha_1 \otimes_{\mathcal{A}} \alpha_2 \vdash \alpha}{a\,\alpha_1 \otimes_{\mathcal{A}} \alpha_2 \vdash a\,\alpha}\ a \notin \mathcal{A} \quad \frac{\alpha_1 \otimes_{\mathcal{A}} \alpha_2 \vdash \alpha}{\alpha_1 \otimes_{\mathcal{A}} a\,\alpha_2 \vdash a\,\alpha}\ a \notin \mathcal{A}$$

Moreover, for $X, Y \subseteq A^* \times \mathbb{R}_{]0,1]}$ we let:

$$X \otimes_{\mathcal{A}} Y \;=\; \{(\alpha, p_1 \cdot p_2) \mid (\alpha_1, p_1) \in X \wedge (\alpha_2, p_2) \in Y \wedge \alpha_1 \otimes_{\mathcal{A}} \alpha_2 \vdash \alpha\}$$

In the rest of this proof, we show that $\mathit{traces}(s_k, s_0) = \mathit{traces}(s_k) \otimes_{\mathcal{A}} \mathit{traces}(s_0)$ for $k = 1, 2$. This, together with Thm. 3.8, guarantees that if $s_1 \sim_{\mathrm{PTr}} s_2$ then $(s_1, s_0) \sim_{\mathrm{PTr}} (s_2, s_0)$. Indeed, if $s_1 \sim_{\mathrm{PTr}} s_2$, then $\mathit{traces}(s_1) = \mathit{traces}(s_2)$ by Thm. 3.8. Thus:

$$\mathit{traces}(s_1, s_0) \;=\; \mathit{traces}(s_1) \otimes_{\mathcal{A}} \mathit{traces}(s_0) \;=\; \mathit{traces}(s_2) \otimes_{\mathcal{A}} \mathit{traces}(s_0) \;=\; \mathit{traces}(s_2, s_0)$$

and hence $(s_1, s_0) \sim_{\mathrm{PTr}} (s_2, s_0)$ by Thm. 3.8.

To be precise, we show that for all $s \in S_1 \cup S_2$, $s_0 \in S_0$, $\alpha \in A^*$, and $p \in \mathbb{R}_{]0,1]}$ there exists $i \in \mathbb{N}$ such that $(\alpha, p) \in \mathit{traces}_i(s, s_0)$ iff there exist $j, h \le i$ such that $(\alpha, p) \in \mathit{traces}_j(s) \otimes_{\mathcal{A}} \mathit{traces}_h(s_0)$ by proceeding by induction on the length of α.

Base of induction: Let $|\alpha| = 0$, i.e., $\alpha = \varepsilon$. In this case, the result follows directly from the fact that:

$$\mathit{traces}_0(s, s_0) \;=\; \{(\varepsilon, 1)\} \;=\; \{(\varepsilon, 1)\} \otimes_{\mathcal{A}} \{(\varepsilon, 1)\} \;=\; \mathit{traces}_0(s) \otimes_{\mathcal{A}} \mathit{traces}_0(s_0)$$

Induction hypothesis: We assume that for all $s' \in S_1 \cup S_2$, $s'_0 \in S_0$, $\alpha' \in A^*$, and $p' \in \mathbb{R}_{]0,1]}$ there exists $i' \in \mathbb{N}$ such that $(\alpha', p') \in \mathit{traces}_{i'}(s', s'_0)$ iff there exist $j', h' \le i'$ such that $(\alpha', p') \in \mathit{traces}_{j'}(s') \otimes_{\mathcal{A}} \mathit{traces}_{h'}(s'_0)$ when $|\alpha| \le n$ for some $n \in \mathbb{N}$.

Induction step: Let $\alpha = a\,\alpha'$ with $|\alpha'| = n$. The fact that $(\alpha, p) \in \mathit{traces}_{n+1}(s, s_0)$ means that there exists a transition $(s, s_0) \xrightarrow{a} \mathcal{D}$ such that:

$$(\alpha', p) \;\in \sum_{(s', s'_0) \in \mathit{supp}(\mathcal{D})} \mathcal{D}(s', s'_0) \cdot \mathit{traces}_n(s', s'_0)$$

where $\longrightarrow$ is the transition relation of $\mathcal{L}_1 \parallel_{\mathcal{A}} \mathcal{L}_0$ or $\mathcal{L}_2 \parallel_{\mathcal{A}} \mathcal{L}_0$ depending on whether s belongs to S_1 or S_2. Similarly, we denote by $\longrightarrow_{1,2}$ the transition relation of $\mathcal{L}_1$ or $\mathcal{L}_2$ depending on whether s belongs to S_1 or S_2.

We distinguish two cases: $a \in \mathcal{A}$ and $a \notin \mathcal{A}$. If $a \in \mathcal{A}$, then $(s, s_0) \xrightarrow{a} \mathcal{D}$ means that $s \xrightarrow{a}_{1,2} \mathcal{D}'$, $s_0 \xrightarrow{a}_0 \mathcal{D}''$, and $\mathcal{D}(s', s'_0) = \mathcal{D}'(s') \cdot \mathcal{D}''(s'_0)$ for all $(s', s'_0) \in (S_1 \cup S_2) \times S_0$, hence:

$$(\alpha', p) \;\in \sum_{s' \in \mathit{supp}(\mathcal{D}')} \; \sum_{s'_0 \in \mathit{supp}(\mathcal{D}'')} \mathcal{D}'(s') \cdot \mathcal{D}''(s'_0) \cdot \mathit{traces}_n(s', s'_0)$$

This means that for each $s' \in \mathit{supp}(\mathcal{D}')$ and $s'_0 \in \mathit{supp}(\mathcal{D}'')$ there exists $p_{(s', s'_0)} \in \mathbb{R}_{[0,1]}$ such that $\mathit{traces}_n(s', s'_0) \vdash (\alpha', p_{(s', s'_0)})$, and $p = \sum_{s' \in \mathit{supp}(\mathcal{D}')} \sum_{s'_0 \in \mathit{supp}(\mathcal{D}'')} \mathcal{D}'(s') \cdot \mathcal{D}''(s'_0) \cdot p_{(s', s'_0)}$.

By applying the induction hypothesis to all $s' \in \mathit{supp}(\mathcal{D}')$ and $s'_0 \in \mathit{supp}(\mathcal{D}'')$ such that $(\alpha', p_{(s', s'_0)}) \in \mathit{traces}_n(s', s'_0)$ and exploiting Lemma 3.6 so as to obtain a single pair from the various pairs $j_{(s', s'_0)}, h_{(s', s'_0)} \le n$, it follows that the fact that $(\alpha, p) \in \mathit{traces}_{n+1}(s, s_0)$ means that there exist $j, h \leq n$ such that:

$$\begin{aligned}
(\alpha, p) \ &\in \ a.(\sum_{s' \in supp(\mathcal{D}')} \sum_{s'_0 \in supp(\mathcal{D}'')} \mathcal{D}'(s') \cdot \mathcal{D}''(s'_0) \cdot (traces_j(s') \otimes_{\mathcal{A}} traces_h(s'_0))) \\
&= \ a.(\sum_{s' \in supp(\mathcal{D}')} \sum_{s'_0 \in supp(\mathcal{D}'')} (\mathcal{D}'(s') \cdot traces_j(s')) \otimes_{\mathcal{A}} (\mathcal{D}''(s'_0) \cdot traces_h(s'_0))) \\
&= \ a.((\sum_{s' \in supp(\mathcal{D}')} \mathcal{D}'(s') \cdot traces_j(s')) \otimes_{\mathcal{A}} (\sum_{s'_0 \in supp(\mathcal{D}'')} \mathcal{D}''(s'_0) \cdot traces_h(s'_0))) \\
&= \ a.(\sum_{s' \in supp(\mathcal{D}')} \mathcal{D}'(s') \cdot traces_j(s')) \otimes_{\mathcal{A}} a.(\sum_{s'_0 \in supp(\mathcal{D}'')} \mathcal{D}''(s'_0) \cdot traces_h(s'_0)) \\
&\subseteq \ traces_{j+1}(s) \otimes_{\mathcal{A}} traces_{h+1}(s_0)
\end{aligned}$$

Similarly, if $a \notin \mathcal{A}$, then $(s, s_0) \xrightarrow{a} \mathcal{D}$ means that either $s \xrightarrow{a}_{1,2} \mathcal{D}'$ with $\mathcal{D}(s', s'_0) = \mathcal{D}'(s')$ if $s'_0 = s_0$ and $\mathcal{D}(s', s'_0) = 0$ if $s'_0 \in S_0 \setminus \{s_0\}$, or $s_0 \xrightarrow{a}_0 \mathcal{D}''$ with $\mathcal{D}(s', s'_0) = \mathcal{D}''(s'_0)$ if $s' = s$ and $\mathcal{D}(s', s'_0) = 0$ if $s' \in (S_1 \cup S_2) \setminus \{s\}$, hence:

$$(\alpha', p) \ \in \sum_{s' \in supp(\mathcal{D}')} \mathcal{D}'(s') \cdot traces_n(s', s_0) \cup \sum_{s'_0 \in supp(\mathcal{D}'')} \mathcal{D}''(s'_0) \cdot traces_n(s, s'_0)$$

This means that (i) for each $s' \in supp(\mathcal{D}')$ there exists $p_{s'} \in \mathbb{R}_{[0,1]}$ such that $traces_n(s', s_0) \vdash (\alpha', p_{s'})$, (ii) for each $s'_0 \in supp(\mathcal{D}'')$ there exists $p_{s'_0} \in \mathbb{R}_{[0,1]}$ such that $traces_n(s, s'_0) \vdash (\alpha', p_{s'_0})$, and (iii) either $p = \sum_{s' \in supp(\mathcal{D}')} \mathcal{D}'(s') \cdot p_{s'}$ or $p = \sum_{s'_0 \in supp(\mathcal{D}'')} \mathcal{D}''(s'_0) \cdot p_{s'_0}$. By applying the induction hypothesis to all $s' \in supp(\mathcal{D}')$ such that $(\alpha', p_{s'}) \in traces_n(s', s_0)$ and to all $s'_0 \in supp(\mathcal{D}'')$ such that $(\alpha', p_{s'_0}) \in traces_n(s, s'_0)$, and exploiting Lemma 3.6 so as to obtain a single pair from the various pairs $j_{s'}, h_{s'} \leq n$ and $j_{s'_0}, h_{s'_0} \leq n$, it follows that the fact that $(\alpha, p) \in traces_{n+1}(s, s_0)$ means that there exist $j, h \leq n$ such that:

$$\begin{aligned}
(\alpha, p) \ &\in \ a.(\sum_{s' \in supp(\mathcal{D}')} \mathcal{D}'(s') \cdot (traces_j(s') \otimes_{\mathcal{A}} traces_h(s_0))) \cup \\
&\quad\ a.(\sum_{s'_0 \in supp(\mathcal{D}'')} \mathcal{D}''(s'_0) \cdot (traces_j(s) \otimes_{\mathcal{A}} traces_h(s'_0))) \\
&= \ a.((\sum_{s' \in supp(\mathcal{D}')} \mathcal{D}'(s') \cdot traces_j(s')) \otimes_{\mathcal{A}} traces_h(s_0)) \cup \\
&\quad\ a.(traces_j(s) \otimes_{\mathcal{A}} (\sum_{s'_0 \in supp(\mathcal{D}'')} \mathcal{D}''(s'_0) \cdot traces_h(s'_0))) \\
&= \ (a.(\sum_{s' \in supp(\mathcal{D}')} \mathcal{D}'(s') \cdot traces_j(s'))) \otimes_{\mathcal{A}} traces_h(s_0) \cup \\
&\quad\ traces_j(s) \otimes_{\mathcal{A}} (a.(\sum_{s'_0 \in supp(\mathcal{D}'')} \mathcal{D}''(s'_0) \cdot traces_h(s'_0))) \\
&\subseteq \ traces_{j+1}(s) \otimes_{\mathcal{A}} traces_h(s_0) \cup \\
&\quad\ traces_j(s) \otimes_{\mathcal{A}} traces_{h+1}(s_0) \\
&\subseteq \ traces_{j+1}(s) \otimes_{\mathcal{A}} traces_{h+1}(s_0) \cup \\
&\quad\ traces_{j+1}(s) \otimes_{\mathcal{A}} traces_{h+1}(s_0) \\
&= \ traces_{j+1}(s) \otimes_{\mathcal{A}} traces_{h+1}(s_0)
\end{aligned}$$

where we have exploited again Lemma 3.6. □

It can be similarly proved that also $\sim^{\mathrm{ct}}_{\mathrm{PTr}}$ is a congruence with respect to parallel composition if combined transitions are considered instead of ordinary ones in Def. 3.5.

Theorem 3.10. *Let $\mathcal{L}_k = (S_k, A, \longrightarrow_k)$ be an NPLTS for $k = 0, 1, 2$ and consider $\mathcal{L}_1 \,\|_{\mathcal{A}}\, \mathcal{L}_0$ and $\mathcal{L}_2 \,\|_{\mathcal{A}}\, \mathcal{L}_0$ for $\mathcal{A} \subseteq A$. Let $s_k \in S_k$ for $k = 0, 1, 2$. Then:*

$$s_1 \sim^{\mathrm{ct}}_{\mathrm{PTr}} s_2 \implies (s_1, s_0) \sim^{\mathrm{ct}}_{\mathrm{PTr}} (s_2, s_0)$$ □

4. Testing Equivalences for NPLTS Models

Testing equivalences for NPLTS models consider the probability of performing computations along which the same tests are passed. Tests specify the actions a process can perform; in this setting, tests are formalized as NPLTS models equipped with a success state. For the sake of simplicity, we restrict ourselves to *finite* tests, each of which has finitely many states, finitely many outgoing transitions from each state, an acyclic graph structure, and hence finitely many computations leading to the success state.

Definition 4.1. A *nondeterministic and probabilistic test*, NPT for short, is a finite NPLTS $\mathcal{T} = (O, A, \longrightarrow)$ where O contains a distinguished success state denoted by ω with no outgoing transitions. We say that a computation of $\mathcal{T}$ is *successful* iff its last state is ω. ■

Definition 4.2. Let $\mathcal{L} = (S, A, \longrightarrow)$ be an NPLTS and $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ be an NPT. The *interaction system* of $\mathcal{L}$ and $\mathcal{T}$ is the NPLTS $\mathcal{I}(\mathcal{L}, \mathcal{T}) = \mathcal{L} \parallel \mathcal{T}$ where:

- Every element $(s, o) \in S \times O$ is called a *configuration* and is said to be *successful* iff $o = \omega$.
- A computation of $\mathcal{I}(\mathcal{L}, \mathcal{T})$ is said to be *successful* iff its last configuration is successful. Given a resolution $\mathcal{Z}$ of $(s, o) \in S \times O$, we denote by $\mathcal{SC}(z_{s,o})$ the set of successful computations from the state $z_{s,o}$ of $\mathcal{Z}$ corresponding to the configuration (s, o) of $\mathcal{I}(\mathcal{L}, \mathcal{T})$.

■

In the following, we shall consider only maximal resolutions of interaction systems because the non-maximal ones do not expose all successful computations.

Due to the possible presence of equally labeled transitions departing from the same state, there is not necessarily a single probability value with which an NPLTS passes a test. Thus, to compare two states s_1 and s_2 of an NPLTS via an NPT with initial state o, we need to compute the probability of performing a successful computation from the two configurations (s_1, o) and (s_2, o) in every maximal resolution of the interaction system. As done in [39, 21, 33, 12], one option is comparing only the suprema ($\sqcup$) and the infima ($\sqcap$) of these success probabilities over all maximal resolutions of the interaction systems. To avoid infima to be trivially zero, it is strictly necessary to consider only maximal resolutions.

Definition 4.3. Let $(S, A, \longrightarrow)$ be an NPLTS. We say that $s_1, s_2 \in S$ are *probabilistic* $\sqcup\sqcap$*-testing equivalent*, written $s_1 \sim_{\text{PTe-}\sqcup\sqcap} s_2$, iff for every NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$ it holds that:

$$\bigsqcup_{\mathcal{Z}_1 \in Res_{\max}(s_1, o)} prob(\mathcal{SC}(z_{s_1,o})) = \bigsqcup_{\mathcal{Z}_2 \in Res_{\max}(s_2, o)} prob(\mathcal{SC}(z_{s_2,o}))$$
$$\bigsqcap_{\mathcal{Z}_1 \in Res_{\max}(s_1, o)} prob(\mathcal{SC}(z_{s_1,o})) = \bigsqcap_{\mathcal{Z}_2 \in Res_{\max}(s_2, o)} prob(\mathcal{SC}(z_{s_2,o}))$$

We denote by $\sim^{\text{ct}}_{\text{PTe-}\sqcup\sqcap}$ the variant based on randomized schedulers. ■

Following the structure of classical testing equivalence $\sim_{\text{Te,fnd}}$ for fully nondeterministic processes [11], the constraint on suprema represents the may-part of $\sim_{\text{PTe-}\sqcup\sqcap}$ while the constraint on infima represents the must-part of $\sim_{\text{PTe-}\sqcup\sqcap}$. The probabilistic testing equivalences in [39, 21, 12] are essentially defined as $\sim_{\text{PTe-}\sqcup\sqcap}$, while the one in [33] resolves nondeterminism through randomized schedulers instead of deterministic ones and makes use of countably many success actions in place of a single one. Notably, a single success action suffices when testing finitary processes, as proved in [14], and the use of different classes of schedulers does not change the discriminating power, as we now show.

Theorem 4.4. *Let* $(S, A, \longrightarrow)$ *be an NPLTS and* $s_1, s_2 \in S$. *Then:*

$$s_1 \sim^{\text{ct}}_{\text{PTe-}\sqcup\sqcap} s_2 \iff s_1 \sim_{\text{PTe-}\sqcup\sqcap} s_2$$

Proof. The result follows from the fact that, given an arbitrary state $s \in S$ and an arbitrary NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$, it holds that:

$$\bigsqcup_{\mathcal{Z} \in Res^{\rm ct}_{\max}(s,o)} prob(\mathcal{SC}(z_{s,o})) \; = \bigsqcup_{\mathcal{Z} \in Res_{\max}(s,o)} prob(\mathcal{SC}(z_{s,o}))$$

and an analogous equality holds for infima. In fact, first of all we note that:

$$\bigsqcup_{\mathcal{Z} \in Res^{\rm ct}_{\max}(s,o)} prob(\mathcal{SC}(z_{s,o})) \; \geq \bigsqcup_{\mathcal{Z} \in Res_{\max}(s,o)} prob(\mathcal{SC}(z_{s,o}))$$

because a deterministic scheduler is a special case of randomized scheduler and hence the set of probabilities on the left contains the set of probabilities on the right (a dual property based on $\leq$ holds for infima). Therefore, it suffices to show that:

$$\bigsqcup_{\mathcal{Z} \in Res^{\rm ct}_{\max}(s,o)} prob(\mathcal{SC}(z_{s,o})) \; \leq \bigsqcup_{\mathcal{Z} \in Res_{\max}(s,o)} prob(\mathcal{SC}(z_{s,o}))$$

which we prove below by proceeding by induction on the length n of the longest successful computation from (s, o), which is finite because $\mathcal{T}$ is finite (a dual property based on $\geq$ can be established for infima):

- If $n = 0$, i.e., $o = \omega$, then:
$$\bigsqcup_{\mathcal{Z} \in Res^{\rm ct}_{\max}(s,o)} prob(\mathcal{SC}(z_{s,o})) \; = \; 1 \; = \bigsqcup_{\mathcal{Z} \in Res_{\max}(s,o)} prob(\mathcal{SC}(z_{s,o}))$$
- Let $n \in \mathbb{N}_{>0}$ and suppose that the property holds for all configurations from which the longest successful computation has length $m = 0, \dots, n-1$. Indicating with $(s, o) \xrightarrow{a}_{\rm c} \mathcal{D}_{\rm c}$ a combined transition from (s, o) with $\mathcal{D}_{\rm c} = \sum_{i=1}^{m} p_i \cdot \mathcal{D}_i$, we have that:
$$\begin{aligned}
&\bigsqcup_{\mathcal{Z} \in Res^{\rm ct}_{\max}(s,o)} prob(\mathcal{SC}(z_{s,o})) \; = \\
&= \bigsqcup_{(s,o) \xrightarrow{a}_{\rm c} \mathcal{D}_{\rm c}} \sum_{(s',o') \in S \times O} \left(\mathcal{D}_{\rm c}(s', o') \cdot \bigsqcup_{\mathcal{Z}' \in Res^{\rm ct}_{\max}(s',o')} prob(\mathcal{SC}(z'_{s',o'})) \right) \\
&\leq \bigsqcup_{(s,o) \xrightarrow{a}_{\rm c} \mathcal{D}_{\rm c}} \sum_{(s',o') \in S \times O} \left(\mathcal{D}_{\rm c}(s', o') \cdot \bigsqcup_{\mathcal{Z}' \in Res_{\max}(s',o')} prob(\mathcal{SC}(z'_{s',o'})) \right) \\
&= \bigsqcup_{(s,o) \xrightarrow{a}_{\rm c} \mathcal{D}_{\rm c}} \sum_{(s',o') \in S \times O} \left(\sum_{i=1}^{m} (p_i \cdot \mathcal{D}_i(s', o')) \cdot \bigsqcup_{\mathcal{Z}' \in Res_{\max}(s',o')} prob(\mathcal{SC}(z'_{s',o'})) \right) \\
&= \bigsqcup_{(s,o) \xrightarrow{a}_{\rm c} \mathcal{D}_{\rm c}} \sum_{i=1}^{m} p_i \cdot \left(\sum_{(s',o') \in S \times O} \left(\mathcal{D}_i(s', o') \cdot \bigsqcup_{\mathcal{Z}' \in Res_{\max}(s',o')} prob(\mathcal{SC}(z'_{s',o'})) \right) \right) \\
&\leq \bigsqcup_{(s,o) \xrightarrow{a}_{\rm c} \mathcal{D}_{\rm c}} \sum_{i=1}^{m} p_i \cdot \bigsqcup_{i=1}^{m} \left(\sum_{(s',o') \in S \times O} \left(\mathcal{D}_i(s', o') \cdot \bigsqcup_{\mathcal{Z}' \in Res_{\max}(s',o')} prob(\mathcal{SC}(z'_{s',o'})) \right) \right) \\
&= \bigsqcup_{(s,o) \xrightarrow{a}_{\rm c} \mathcal{D}_{\rm c}} \bigsqcup_{i=1}^{m} \left(\sum_{(s',o') \in S \times O} \left(\mathcal{D}_i(s', o') \cdot \bigsqcup_{\mathcal{Z}' \in Res_{\max}(s',o')} prob(\mathcal{SC}(z'_{s',o'})) \right) \right) \\
&= \bigsqcup_{(s,o) \xrightarrow{a} \mathcal{D}} \sum_{(s',o') \in S \times O} \left(\mathcal{D}(s', o') \cdot \bigsqcup_{\mathcal{Z}' \in Res_{\max}(s',o')} prob(\mathcal{SC}(z'_{s',o'})) \right) \\
&= \bigsqcup_{\mathcal{Z} \in Res_{\max}(s,o)} prob(\mathcal{SC}(z_{s,o}))
\end{aligned}$$
where in the third line we have exploited the induction hypothesis and in the seventh line the fact that $\sum_{i=1}^{m} p_i = 1$. ∎

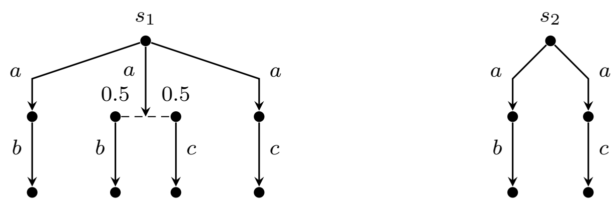

Figure 5: NPLTS models identified by $\sim^{\rm ct}_{\rm PTr,dis} / \sim_{\rm PTe\text{-}\sqcup\sqcap}$ and told apart by $\sim_{\rm PTr,dis} / \sim_{\rm PTr}$

The relation $\sim_{\rm PTe\text{-}\sqcup\sqcap}$ does not enjoy the desirable property – possessed by $\sim_{\rm Te,fnd}$ – of resulting in a testing semantics finer than the trace semantics for the same class of processes. Whether $\sim_{\rm PTe\text{-}\sqcup\sqcap}$ is included in the trace equivalences of Sect. 3 depends on the type of schedulers that are considered on the trace semantics side. In the case of randomized schedulers, as shown in [33] it holds that $\sim_{\rm PTe\text{-}\sqcup\sqcap} \subseteq \sim^{\rm ct}_{\rm PTr,dis}$, and hence $\sim_{\rm PTe\text{-}\sqcup\sqcap} \subseteq \sim^{\rm ct}_{\rm PTr}$ by virtue of Thm. 3.3. However, inclusion no longer holds when only deterministic schedulers are admitted. Let us consider the two NPLTS models in Fig. 5. We have that $s_1 \sim_{\rm PTe\text{-}\sqcup\sqcap} s_2$ while $s_1 \not\sim_{\rm PTr,dis} s_2$ and $s_1 \not\sim_{\rm PTr} s_2$. It holds that $s_1 \sim_{\rm PTe\text{-}\sqcup\sqcap} s_2$ because, for any test, the central maximal resolution of s_1 always gives rise to a success probability comprised between the success probabilities of the other two maximal resolutions of s_1, which correspond to the two maximal resolutions of s_2. In contrast, s_1 and s_2 are not related by the two probabilistic trace equivalences because the maximal resolution of s_1 starting with the central a-transition is not matched by any of the two maximal resolutions of s_2.

Under deterministic schedulers, inclusion can be achieved by considering $\sim_{\rm PTr}$ in lieu of the finer $\sim_{\rm PTr,dis}$ and the new testing equivalence $\sim_{\rm PTe\text{-}\forall\exists}$ introduced by the next definition in lieu of the coarser $\sim_{\rm PTe\text{-}\sqcup\sqcap}$. Instead of focussing only on extremal success probabilities, $\sim_{\rm PTe\text{-}\forall\exists}$ requires matching the success probabilities of *all* maximal resolutions of the interaction systems. Interestingly, the variant of $\sim_{\rm PTe\text{-}\forall\exists}$ based on randomized schedulers coincides with $\sim_{\rm PTe\text{-}\sqcup\sqcap}$.

Definition 4.5. Let $(S, A, \longrightarrow)$ be an NPLTS. We say that $s_1, s_2 \in S$ are *probabilistic $\forall\exists$-testing equivalent*, written $s_1 \sim_{\rm PTe\text{-}\forall\exists} s_2$, iff for every NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$ it holds that:

- For each $\mathcal{Z}_1 \in Res_{\max}(s_1, o)$ there exists $\mathcal{Z}_2 \in Res_{\max}(s_2, o)$ such that:
$$prob(\mathcal{SC}(z_{s_1,o})) \; = \; prob(\mathcal{SC}(z_{s_2,o}))$$
- For each $\mathcal{Z}_2 \in Res_{\max}(s_2, o)$ there exists $\mathcal{Z}_1 \in Res_{\max}(s_1, o)$ such that:
$$prob(\mathcal{SC}(z_{s_2,o})) \; = \; prob(\mathcal{SC}(z_{s_1,o}))$$

We denote by $\sim^{\rm ct}_{\rm PTe\text{-}\forall\exists}$ the coarser variant based on randomized schedulers. ■

Theorem 4.6. *Let $(S, A, \longrightarrow)$ be an image-finite NPLTS and $s_1, s_2 \in S$. Then:*
$$s_1 \sim_{\rm PTe\text{-}\forall\exists} s_2 \implies s_1 \sim_{\rm PTe\text{-}\sqcup\sqcap} s_2$$
$$s_1 \sim^{\rm ct}_{\rm PTe\text{-}\forall\exists} s_2 \iff s_1 \sim_{\rm PTe\text{-}\sqcup\sqcap} s_2$$

Proof. If $s_1 \sim_{\rm PTe\text{-}\forall\exists} s_2$, then we immediately derive that for every NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$ it holds that:
$$\{prob(\mathcal{SC}(z_{s_1,o})) \mid \mathcal{Z}_1 \in Res_{\max}(s_1, o)\} \subseteq \{prob(\mathcal{SC}(z_{s_2,o})) \mid \mathcal{Z}_2 \in Res_{\max}(s_2, o)\}$$
$$\{prob(\mathcal{SC}(z_{s_2,o})) \mid \mathcal{Z}_2 \in Res_{\max}(s_2, o)\} \subseteq \{prob(\mathcal{SC}(z_{s_1,o})) \mid \mathcal{Z}_1 \in Res_{\max}(s_1, o)\}$$
As a consequence:
$$\{prob(\mathcal{SC}(z_{s_1,o})) \mid \mathcal{Z}_1 \in Res_{\max}(s_1, o)\} \; = \; \{prob(\mathcal{SC}(z_{s_2,o})) \mid \mathcal{Z}_2 \in Res_{\max}(s_2, o)\}$$

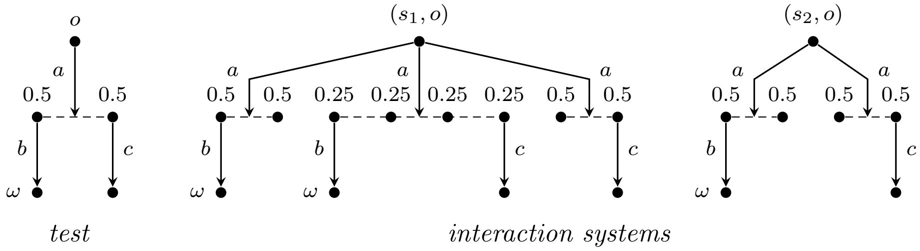

Figure 6: A test showing that the two NPLTS models in Fig. 5 are distinguished by $\sim_{\text{PTe-}\forall\exists}$

and hence:

$$\begin{aligned}
\bigsqcup_{\mathcal{Z}_1 \in Res_{\max}(s_1,o)} prob(\mathcal{SC}(z_{s_1,o})) &= \bigsqcup_{\mathcal{Z}_2 \in Res_{\max}(s_2,o)} prob(\mathcal{SC}(z_{s_2,o})) \\
\bigsqcap_{\mathcal{Z}_1 \in Res_{\max}(s_1,o)} prob(\mathcal{SC}(z_{s_1,o})) &= \bigsqcap_{\mathcal{Z}_2 \in Res_{\max}(s_2,o)} prob(\mathcal{SC}(z_{s_2,o}))
\end{aligned}$$

which means that $s_1 \sim_{\text{PTe-}\sqcup\sqcap} s_2$.

The fact that $s_1 \sim^{\text{ct}}_{\text{PTe-}\forall\exists} s_2 \implies s_1 \sim_{\text{PTe-}\sqcup\sqcap} s_2$ stems from $s_1 \sim^{\text{ct}}_{\text{PTe-}\forall\exists} s_2 \implies s_1 \sim^{\text{ct}}_{\text{PTe-}\sqcup\sqcap} s_2$ (as a consequence of the previous result) and $\sim^{\text{ct}}_{\text{PTe-}\sqcup\sqcap} = \sim_{\text{PTe-}\sqcup\sqcap}$ (by virtue of Thm. 4.4).

Suppose now that $s_1 \sim_{\text{PTe-}\sqcup\sqcap} s_2$ and consider an arbitrary NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$, so that:

$$\begin{aligned}
\bigsqcup_{\mathcal{Z}_1 \in Res_{\max}(s_1,o)} prob(\mathcal{SC}(z_{s_1,o})) &= p_\sqcup = \bigsqcup_{\mathcal{Z}_2 \in Res_{\max}(s_2,o)} prob(\mathcal{SC}(z_{s_2,o})) \\
\bigsqcap_{\mathcal{Z}_1 \in Res_{\max}(s_1,o)} prob(\mathcal{SC}(z_{s_1,o})) &= p_\sqcap = \bigsqcap_{\mathcal{Z}_2 \in Res_{\max}(s_2,o)} prob(\mathcal{SC}(z_{s_2,o}))
\end{aligned}$$

If $p_\sqcup = p_\sqcap$, then all the maximal resolutions of (s_1, o) and (s_2, o) have the same success probability, from which it trivially follows that $s_1 \sim_{\text{PTe-}\forall\exists} s_2$ and hence $s_1 \sim^{\text{ct}}_{\text{PTe-}\forall\exists} s_2$.

Recalling that the NPLTS is image finite and the test is finite so that $Res_{\max}(s_1, o)$ and $Res_{\max}(s_2, o)$ are both finite, if $p_\sqcup > p_\sqcap$, then $p_\sqcup$ must be achieved on $\mathcal{Z}_{1,\sqcup} \in Res_{\max}(s_1, o)$ and $\mathcal{Z}_{2,\sqcup} \in Res_{\max}(s_2, o)$ exhibiting the same successful traces, otherwise – observing that both resolutions must have at least one successful trace, otherwise it would be $p_\sqcup = 0$ thus violating $p_\sqcup > p_\sqcap$ – states s_1 and s_2 would be distinguished with respect to $\sim_{\text{PTe-}\sqcup\sqcap}$ by a test obtained from $\mathcal{T}$ by making success reachable only along the successful traces of the one of $\mathcal{Z}_{1,\sqcup}$ and $\mathcal{Z}_{2,\sqcup}$ having a successful trace not possessed by the other, unless that resolution also contains all the successful traces of the other resolution, in which case success must be made reachable only along the successful traces of the other resolution in order to contradict $s_1 \sim_{\text{PTe-}\sqcup\sqcap} s_2$.

Likewise, $p_\sqcap$ must be achieved on $\mathcal{Z}_{1,\sqcap} \in Res_{\max}(s_1, o)$ and $\mathcal{Z}_{2,\sqcap} \in Res_{\max}(s_2, o)$ exhibiting the same unsuccessful maximal traces, otherwise – observing that both resolutions must have at least one unsuccessful maximal trace, otherwise it would be $p_\sqcap = 1$ thus violating $p_\sqcup > p_\sqcap$ – states s_1 and s_2 would be distinguished with respect to $\sim_{\text{PTe-}\sqcup\sqcap}$ by a test obtained from $\mathcal{T}$ by making success reachable also along an unsuccessful maximal trace occurring only in either $\mathcal{Z}_{1,\sqcap}$ or $\mathcal{Z}_{2,\sqcap}$.

By reasoning on the dual test $\mathcal{T}'$ in which the final states of $\mathcal{T}$ that are successful (resp. unsuccessful) are made unsuccessful (resp. successful), it turns out that $\mathcal{Z}_{1,\sqcup}$ and $\mathcal{Z}_{2,\sqcup}$ must also exhibit the same unsuccessful maximal traces and that $\mathcal{Z}_{1,\sqcap}$ and $\mathcal{Z}_{2,\sqcap}$ must also exhibit the same successful traces.

If $\mathcal{Z}_{1,\sqcup}$ and $\mathcal{Z}_{2,\sqcup}$ do not have sequences of initial transitions in common with $\mathcal{Z}_{1,\sqcap}$ and $\mathcal{Z}_{2,\sqcap}$,

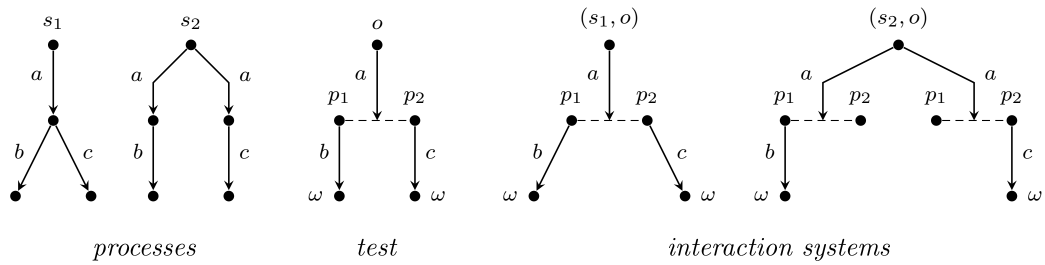

Figure 7: NPLTS models identified by $\sim_{\rm PTr}$ and told apart by $\sim_{\text{PTe-}\forall\exists}$

then $\mathcal{Z}_{1,\sqcup}$ and $\mathcal{Z}_{1,\sqcap}$ on one side and $\mathcal{Z}_{2,\sqcup}$ and $\mathcal{Z}_{2,\sqcap}$ on the other side cannot generate via convex combinations any new resolution that would arise from a randomized scheduler, otherwise they can generate all such resolutions having a certain sequence of initial transitions, thus covering all the intermediate success probabilities between $p_\sqcup$ and $p_\sqcap$ for that sequence of initial transitions. This shows that for each $\mathcal{Z}_1 \in Res^{\rm ct}_{\max}(s_1, o)$ with that sequence of initial transitions there exists $\mathcal{Z}_2 \in Res^{\rm ct}_{\max}(s_2, o)$ with that sequence of initial transitions such that $prob(\mathcal{SC}(z_{s_1,o})) = prob(\mathcal{SC}(z_{s_2,o}))$, and vice versa.
The same procedure can now be applied to the remaining resolutions in $Res_{\max}(s_1, o)$ and $Res_{\max}(s_2, o)$ that are not convex combinations of previously considered resolutions, starting from those among the remaining resolutions on which the maximal and minimal success probabilities are achieved. We can thus conclude that $s_1 \sim^{\rm ct}_{\text{PTe-}\forall\exists} s_2$. □

The inclusion of $\sim_{\text{PTe-}\forall\exists}$ in $\sim_{\text{PTe-}\sqcup\sqcap}$ is strict. Indeed, if we consider again the two $\sim_{\text{PTe-}\sqcup\sqcap}$-equivalent NPLTS models in Fig. 5 and we apply the test in Fig. 6, it turns out that $s_1 \not\sim_{\text{PTe-}\forall\exists} s_2$. For the two interaction systems in Fig. 6, we have that the maximal resolution of (s_1, o) starting with the central a-transition gives rise to a success probability equal to 0.25 that is not matched by any of the two maximal resolutions of (s_2, o). These resolutions, which correspond to the maximal resolutions of (s_1, o) starting with the two outermost a-transitions, have success probability 0.5 and 0, respectively.

Theorem 4.7. *Let $(S, A, \longrightarrow)$ be an NPLTS and $s_1, s_2 \in S$. Then:*
$$s_1 \sim_{\text{PTe-}\forall\exists} s_2 \implies s_1 \sim_{\rm PTr} s_2$$

Proof. If $s_1 \sim_{\text{PTe-}\forall\exists} s_2$, then in particular for every NPT $\mathcal{T}_\alpha = (O, A, \longrightarrow_{\mathcal{T}_\alpha})$ with initial state $o \in O$ having a single maximal computation that is labeled with $\alpha \in A^*$ and reaches success, it holds that:

- For each $\mathcal{Z}_1 \in Res_{\max}(s_1, o)$ there exists $\mathcal{Z}_2 \in Res_{\max}(s_2, o)$ such that:
$$prob(\mathcal{SC}(z_{s_1,o})) \;=\; prob(\mathcal{SC}(z_{s_2,o}))$$
- For each $\mathcal{Z}_2 \in Res_{\max}(s_2, o)$ there exists $\mathcal{Z}_1 \in Res_{\max}(s_1, o)$ such that:
$$prob(\mathcal{SC}(z_{s_2,o})) \;=\; prob(\mathcal{SC}(z_{s_1,o}))$$

Since $prob(\mathcal{SC}^{\mathcal{Z}}(z_{s,o})) = prob(\mathcal{CC}^{\mathcal{Z}'}(z_s, \alpha))$ for all $s \in S$ due to the structure of $\mathcal{T}_\alpha$ – where $\mathcal{Z} \in Res_{\max}(s, o)$ and $\mathcal{Z}' \in Res(s)$ originates $\mathcal{Z}$ in the interaction with $\mathcal{T}_\alpha$ – we immediately derive that for all $\alpha \in A^*$ it holds that:

- For each $\mathcal{Z}_1 \in Res(s_1)$ there exists $\mathcal{Z}_2 \in Res(s_2)$ such that:
$$prob(\mathcal{CC}(z_{s_1}, \alpha)) \;=\; prob(\mathcal{CC}(z_{s_2}, \alpha))$$

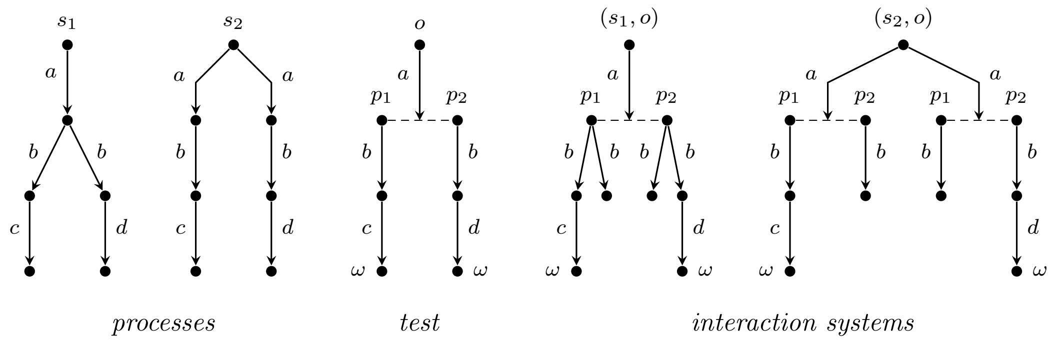

Figure 8: NPLTS models equated by $\sim_{\mathrm{Te,fnd}}$ and distinguished by $\sim_{\text{PTe-}\sqcup\sqcap}/\sim_{\text{PTe-}\forall\exists}$

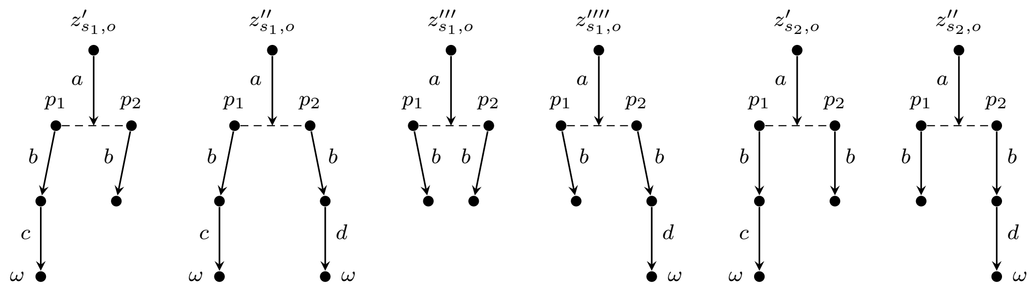

Figure 9: Maximal resolutions of the two interaction systems in Fig. 8.

- For each $\mathcal{Z}_2 \in \mathit{Res}(s_2)$ there exists $\mathcal{Z}_1 \in \mathit{Res}(s_1)$ such that:
$$\mathit{prob}(\mathcal{CC}(z_{s_2}, \alpha)) \;=\; \mathit{prob}(\mathcal{CC}(z_{s_1}, \alpha))$$

This means that $s_1 \sim_{\mathrm{PTr}} s_2$. □

The inclusion of $\sim_{\text{PTe-}\forall\exists}$ in $\sim_{\mathrm{PTr}}$ is strict. For instance, if we consider the two NPLTS models in Fig. 7, it turns out that $s_1 \sim_{\mathrm{PTr}} s_2$ while $s_1 \not\sim_{\text{PTe-}\forall\exists} s_2$. In fact, the test in Fig. 7 distinguishes s_1 from s_2 with respect to $\sim_{\text{PTe-}\forall\exists}$ because – looking at the two interaction systems also reported in the figure – the only maximal resolution of (s_1, o) has a success probability equal to 1 that is not matched by any of the two maximal resolutions of (s_2, o), whose success probabilities are p_1 and p_2, respectively.

Another desirable property of relations like $\sim_{\text{PTe-}\sqcup\sqcap}$ and $\sim_{\text{PTe-}\forall\exists}$ that are defined over a general class of processes is that of being backward compatible with analogous relations for restricted classes of processes. Specifically, we refer to testing equivalences $\sim_{\mathrm{Te,fnd}}$ for fully nondeterministic processes [11], $\sim_{\mathrm{Te,fpr}}$ for fully probabilistic processes [9], and $\sim_{\mathrm{Te,rpr}}$ for reactive probabilistic processes inspired by [25].

As we shall see by means of two counterexamples, backward compatibility is only *partial* as it depends on the set of tests that are used. Intuitively, $\sim_{\text{PTe-}\sqcup\sqcap}$ and $\sim_{\text{PTe-}\forall\exists}$ become sensitive to the moment of occurrence of internal choices when comparing fully nondeterministic processes (resp. fully/reactive probabilistic processes) on the basis of tests admitting probabilities (resp. internal nondeterminism). In such cases, the capability of making copies of intermediate states of the processes under test arises, a fact that in general increases the distinguishing power of testing equivalence, as pointed out in [1]. In a

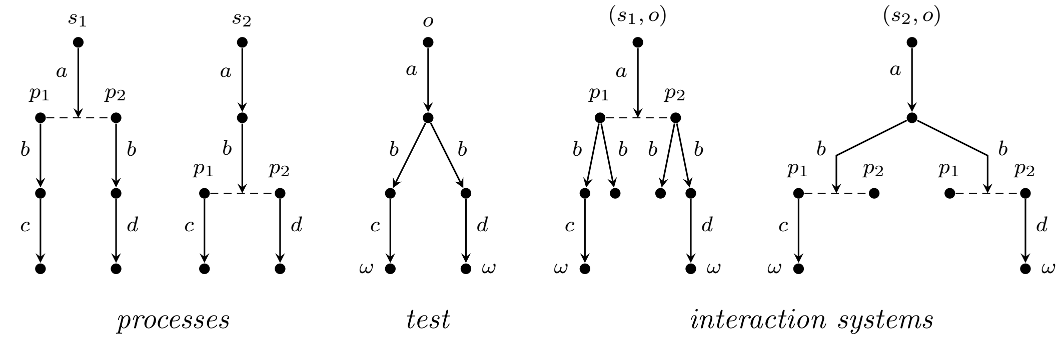

Figure 10: NPLTS models equated by $\sim_{\rm Te,fpr}/\sim_{\rm Te,rpr}$ and distinguished by $\sim_{\text{PTe-}\sqcup\sqcap}/\sim_{\text{PTe-}\forall\exists}$

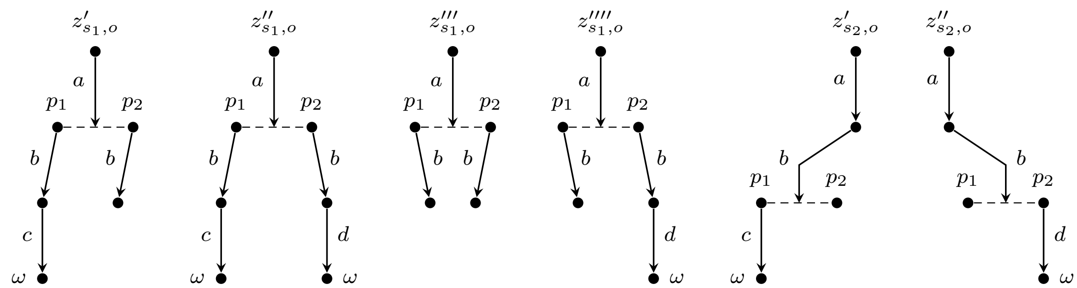

Figure 11: Maximal resolutions of the two interaction systems in Fig. 10.

probabilistic setting, this may lead to questionable estimations of success probabilities (see [16] and the references therein). Indeed, taking advantage of the increased discriminating power, in [12] it was shown that the may-part of $\sim_{\text{PTe-}\sqcup\sqcap}$ coincides with a simulation equivalence akin to the one in [27] and the must-part coincides with a novel failure simulation equivalence. Moreover, in [33] it was shown that the may-part coincides with the coarsest congruence contained in the probabilistic trace-distribution equivalence of [32] and the must-part coincides with the coarsest congruence contained in a probabilistic failure-distribution equivalence.

As observed in [20, 13], it is easy to see that there exist *fully nondeterministic* NPLTS models that are identified by $\sim_{\rm Te,fnd}$ but distinguished by $\sim_{\text{PTe-}\sqcup\sqcap}$ (and also by $\sim_{\text{PTe-}\forall\exists}$). Let us consider the two NPLTS models in Fig. 8, which represent the classical example that illustrates the main difference between testing semantics and bisimulation semantics in a nondeterministic setting. It turns out that $s_1 \sim_{\rm Te,fnd} s_2$ while $s_1 \not\sim_{\text{PTe-}\sqcup\sqcap} s_2$ and $s_1 \not\sim_{\text{PTe-}\forall\exists} s_2$. The *probabilistic* test in Fig. 8 distinguishes s_1 from s_2 with respect to $\sim_{\text{PTe-}\sqcup\sqcap}$. Indeed, if we consider the two interaction systems also reported in Fig. 8 and their maximal resolutions shown in Fig. 9, the supremum of the success probabilities of the four maximal resolutions of (s_1, o) is 1 – see the second maximal resolution of (s_1, o) – whereas the supremum of the success probabilities of the two maximal resolutions of (s_2, o) is equal to the maximum between p_1 and p_2. The same test also distinguishes s_1 from s_2 with respect to $\sim_{\text{PTe-}\forall\exists}$ because the third maximal resolution of (s_1, o) has a success probability equal to 0 that is not matched by any of the two maximal resolutions of (s_2, o), whose success probabilities are p_1 and p_2, respectively.

Following [20], we can easily find also two *fully/reactive probabilistic* NPLTS models that are identified by $\sim_{\mathrm{Te,fpr}}/\sim_{\mathrm{Te,rpr}}$ and distinguished by $\sim_{\mathrm{PTe\text{-}\sqcup\sqcap}}$ (and also by $\sim_{\mathrm{PTe\text{-}\forall\exists}}$). They are depicted in Fig. 10 and constitute the classical example that differentiates probabilistic testing semantics from probabilistic bisimulation semantics. We have that $s_1 \sim_{\mathrm{Te,fpr}} s_2$ and $s_1 \sim_{\mathrm{Te,rpr}} s_2$, while $s_1 \not\sim_{\mathrm{PTe\text{-}\sqcup\sqcap}} s_2$ and $s_1 \not\sim_{\mathrm{PTe\text{-}\forall\exists}} s_2$. The *fully nondeterministic* test in Fig. 10 distinguishes s_1 from s_2 with respect to $\sim_{\mathrm{PTe\text{-}\sqcup\sqcap}}$ and $\sim_{\mathrm{PTe\text{-}\forall\exists}}$, as can be seen from the two interaction systems there reported and their maximal resolutions shown in Fig. 11.

Summing up, the relations $\sim_{\mathrm{PTe\text{-}\sqcup\sqcap}}$ and $\sim_{\mathrm{PTe\text{-}\forall\exists}}$ are backward compatible with respect to testing equivalences defined over restricted classes of processes as long as they only admit tests that belong to the same class as the processes under test.

Theorem 4.8. *Let $\mathcal{L} = (S, A, \longrightarrow)$ be an NPLTS and $s_1, s_2 \in S$.*

(1) *If $\mathcal{L}$ is fully nondeterministic and only fully nondeterministic tests are admitted, then:*
$$s_1 \sim_{\mathrm{PTe\text{-}\sqcup\sqcap}} s_2 \iff s_1 \sim_{\mathrm{PTe\text{-}\forall\exists}} s_2 \iff s_1 \sim_{\mathrm{Te,fnd}} s_2$$
(2) *If $\mathcal{L}$ is fully probabilistic and only fully probabilistic tests are admitted, then:*
$$s_1 \sim_{\mathrm{PTe\text{-}\sqcup\sqcap}} s_2 \iff s_1 \sim_{\mathrm{PTe\text{-}\forall\exists}} s_2 \iff s_1 \sim_{\mathrm{Tr,fpr}} s_2$$
(3) *If $\mathcal{L}$ is reactive probabilistic and only reactive probabilistic tests are admitted, then:*
$$s_1 \sim_{\mathrm{PTe\text{-}\sqcup\sqcap}} s_2 \Longrightarrow s_1 \sim_{\mathrm{Tr,rpr}} s_2$$
$$s_1 \sim_{\mathrm{PTe\text{-}\forall\exists}} s_2 \Longrightarrow s_1 \sim_{\mathrm{Tr,rpr}} s_2$$

Proof. We proceed as follows:

(1) Suppose that $\mathcal{L}$ is fully nondeterministic and that only fully nondeterministic tests are admitted, so that all the resulting interaction systems are fully nondeterministic too. We recall from [11] that $s_1 \sim_{\mathrm{Te,fnd}} s_2$ means that, for every test with initial state o, (i) there exists a successful computation from (s_1, o) iff there exists a successful computation from (s_2, o) and (ii) all maximal computations from (s_1, o) are successful iff all maximal computations from (s_2, o) are successful. The result is a straightforward consequence of the fact that the maximal resolutions of each interaction system coincide with the maximal computations of the interaction system, hence the probability of performing a successful computation within a maximal resolution of an interaction system can only be 1 or 0.

(2) Suppose that $\mathcal{L}$ is fully probabilistic and that only fully probabilistic tests are admitted, so that all the resulting interaction systems are fully probabilistic too. We recall from [9] that $s_1 \sim_{\mathrm{Te,fpr}} s_2$ means that, for every test with initial state o, $prob(\mathcal{SC}(s_1, o)) = prob(\mathcal{SC}(s_2, o))$. The result is a straightforward consequence of the fact that each interaction system has a single maximal resolution, which coincides with the interaction system itself.

(3) Suppose that $\mathcal{L}$ is reactive probabilistic and that only reactive probabilistic tests are admitted, so that all the resulting interaction systems are reactive probabilistic too. Taking inspiration from [25], $s_1 \sim_{\mathrm{Te,rpr}} s_2$ means that, for every test with initial state o, (s_1, o) and (s_2, o) have the same suprema and infima of success probabilities over all of their maximal traces. Success probabilities $prob(\mathcal{SC}_\alpha(s, o))$ are viewed as being conditional on selecting the maximal resolution of (s, o) that contains all the α-compatible computations from (s, o) (this resolution is unique because interaction systems are reactive probabilistic). The result immediately follows by considering tests that reach success along a single trace. □

We conclude with a remark about the four maximal resolutions of (s_1, o) shown in Figs. 9 and 11, whose success probabilities are p_1, 1, 0, and p_2, respectively. The presence of all these resolutions is due to a *demonic* view of nondeterminism, which allows the considered *almighty* schedulers to perform different choices in different copies of the same state of the process under test. This is what happens in the second and in the third maximal resolution, as graphically witnessed by the different orientation of the two b-transitions. In order to be robust with respect to scheduling decisions, these two resolutions cannot be ruled out and their success probabilities, 1 and 0, have to be taken into account.

As pointed out in [6], in a testing scenario schedulers come into play after the process has been composed in parallel with the test, and hence can resolve both local and global nondeterministic choices. This makes it possible for schedulers to make decisions in one component on the basis of the state of the other component, as if there were an information leakage. However, under specific circumstances, one may reasonably consider *less powerful* schedulers ensuring that the choices they perform in different copies of the same state are consistent with each other (see [16] and the references therein). In that case, the two resolutions mentioned above would no longer make sense. As a consequence, values 1 and 0 would respectively become an overestimation and an underestimation of the success probability, and in principle s_1 and s_2 could be identified by $\sim_{\text{PTe-}\sqcup\sqcap}$ and $\sim_{\text{PTe-}\forall\exists}$. We will discuss again the power of schedulers at the end of Sect. 6.

5. Trace-by-Trace Redefinition of Testing Equivalence

In this section, we introduce a new testing equivalence for NPLTS models that is *fully* backward compatible with testing equivalences defined in the literature for restricted classes of processes. In order to counterbalance the stronger discriminating power deriving from the copying capability enabled by tests that do not belong to the class of processes under test, our basic idea is changing the definition of $\sim_{\text{PTe-}\forall\exists}$ by considering success probabilities in a *trace-by-trace fashion* rather than cumulatively over all successful computations of the maximal resolutions.

In the following, given a state s of an NPLTS, a state o of an NPT, and a trace $\alpha \in A^*$, we denote by $Res_{\max,\alpha}(s, o)$ the set of resolutions $\mathcal{Z} \in Res_{\max}(s, o)$ such that $\mathcal{CC}_{\max}(z_{s,o}, \alpha) \neq \emptyset$, where $\mathcal{CC}_{\max}(z_{s,o}, \alpha)$ is the set of computations in $\mathcal{CC}(z_{s,o}, \alpha)$ that are maximal. In other words, $Res_{\max,\alpha}(s, o)$ is the set of maximal resolutions of $z_{s,o}$ having at least one maximal computation labeled with α; the set $Res^{\text{ct}}_{\max,\alpha}(s, o)$ is defined similarly. Moreover, for each resolution $\mathcal{Z}$ we denote by $\mathcal{SCC}(z_{s,o}, \alpha)$ the set of successful α-compatible computations from $z_{s,o}$.

Definition 5.1. Let $(S, A, \longrightarrow)$ be an NPLTS. We say that $s_1, s_2 \in S$ are *probabilistic trace-by-trace testing equivalent*, written $s_1 \sim_{\text{PTe-tbt}} s_2$, iff for every NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$ and <u>for all $\alpha \in A^*$</u> it holds that:

- For each $\mathcal{Z}_1 \in Res_{\max,\alpha}(s_1, o)$ there exists $\mathcal{Z}_2 \in Res_{\max,\alpha}(s_2, o)$ such that:
$$prob(\mathcal{SCC}(z_{s_1,o}, \alpha)) \;=\; prob(\mathcal{SCC}(z_{s_2,o}, \alpha))$$
- For each $\mathcal{Z}_2 \in Res_{\max,\alpha}(s_2, o)$ there exists $\mathcal{Z}_1 \in Res_{\max,\alpha}(s_1, o)$ such that:
$$prob(\mathcal{SCC}(z_{s_2,o}, \alpha)) \;=\; prob(\mathcal{SCC}(z_{s_1,o}, \alpha))$$

We denote by $\sim^{\text{ct}}_{\text{PTe-tbt}}$ the coarser variant based on randomized schedulers. ■

If we consider again the two NPLTS models of Fig. 8 (resp. Fig. 10), it turns out that $s_1 \sim_{\text{PTe-tbt}} s_2$, and hence $s_1 \sim^{\text{ct}}_{\text{PTe-tbt}} s_2$. The interaction of the two processes with the test in the same figure originates maximal computations from (s_1, o) and (s_2, o) that are all labeled with traces $a\,b$, $a\,b\,c$, and $a\,b\,d$. It is easy to see that, in Fig. 9 (resp. Fig. 11), for each of these traces, say α, the probability of performing a successful α-compatible computation in any of the four maximal resolutions of (s_1, o) having a maximal α-compatible computation is matched by the probability of performing a successful α-compatible computation in one of the two maximal resolutions of (s_2, o), and vice versa. As an example, the probability p_1 (resp. p_2) of performing a successful computation compatible with $a\,b\,c$ (resp. $a\,b\,d$) in the second maximal resolution of (s_1, o) is matched by the probability of performing a successful computation compatible with that trace in the first (resp. second) maximal resolution of (s_2, o). As another example, the probability 0 of performing a successful computation compatible with $a\,b$ in the third maximal resolution of (s_1, o) is matched by the probability of performing a successful computation compatible with that trace in any of the two maximal resolutions of (s_2, o).

The examples of Figs. 8 and 10 show that $\sim_{\text{PTe-tbt}}$ and $\sim^{\text{ct}}_{\text{PTe-tbt}}$ are included neither in $\sim_{\text{PTe-}\sqcup\sqcap}$ nor in $\sim_{\text{PTe-}\forall\exists}$. On the other hand, $\sim_{\text{PTe-}\sqcup\sqcap}$ is not included in $\sim_{\text{PTe-tbt}}$ as witnessed by the two NPLTS models in Fig. 5, because the test in Fig. 6 distinguishes s_1 from s_2 with respect to $\sim_{\text{PTe-tbt}}$. In fact, the probability 0.25 of performing a successful computation compatible with $a\,b$ in the maximal resolution of (s_1, o) beginning with the central a-transition is not matched by the probability 0.5 of performing a successful computation compatible with $a\,b$ in the only maximal resolution of (s_2, o) that has a maximal computation labeled with $a\,b$. Thus, $\sim_{\text{PTe-}\sqcup\sqcap}$ and $\sim_{\text{PTe-tbt}}$ are incomparable with each other. What turns out is that $\sim_{\text{PTe-}\sqcup\sqcap}$ is (strictly) included in $\sim^{\text{ct}}_{\text{PTe-tbt}}$, while $\sim_{\text{PTe-}\forall\exists}$ is (strictly) included in $\sim_{\text{PTe-tbt}}$ and hence in $\sim^{\text{ct}}_{\text{PTe-tbt}}$.

Theorem 5.2. *Let $(S, A, \longrightarrow)$ be an NPLTS and $s_1, s_2 \in S$. Then:*

$$s_1 \sim_{\text{PTe-}\sqcup\sqcap} s_2 \implies s_1 \sim^{\text{ct}}_{\text{PTe-tbt}} s_2$$
$$s_1 \sim_{\text{PTe-}\forall\exists} s_2 \implies s_1 \sim_{\text{PTe-tbt}} s_2$$

Proof. Let us initially introduce the following behavioral equivalence: $s_1 \sim_{\text{PTe-tbt},\sqcup\sqcap} s_2$ iff for every NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$ and for all $\alpha \in A^*$ it holds that $Res_{\max,\alpha}(s_1, o) \neq \emptyset$ iff $Res_{\max,\alpha}(s_2, o) \neq \emptyset$ and:

$$\bigsqcup_{\mathcal{Z}_1 \in Res_{\max}(s_1,o)} prob(\mathcal{SCC}(z_{s_1,o}, \alpha)) = \bigsqcup_{\mathcal{Z}_2 \in Res_{\max}(s_2,o)} prob(\mathcal{SCC}(z_{s_2,o}, \alpha))$$
$$\bigsqcap_{\mathcal{Z}_1 \in Res_{\max}(s_1,o)} prob(\mathcal{SCC}(z_{s_1,o}, \alpha)) = \bigsqcap_{\mathcal{Z}_2 \in Res_{\max}(s_2,o)} prob(\mathcal{SCC}(z_{s_2,o}, \alpha))$$

The proof of $s_1 \sim_{\text{PTe-}\sqcup\sqcap} s_2 \implies s_1 \sim^{\text{ct}}_{\text{PTe-tbt}} s_2$ is divided into two parts:

- First, we show that $s_1 \sim_{\text{PTe-}\sqcup\sqcap} s_2 \implies s_1 \sim^{\text{ct}}_{\text{PTe-tbt},\sqcup\sqcap} s_2$. Suppose that $s_1 \sim_{\text{PTe-}\sqcup\sqcap} s_2$ and consider an arbitrary NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$. Given $s \in S$ and $\mathcal{Z} \in Res_{\max}(s, o)$, it holds that:

$$prob(\mathcal{SC}(z_{s,o})) = \sum_{\alpha \in A^* \text{ s.t. } \mathcal{CC}_{\max}(z_{s,o},\alpha) \neq \emptyset} prob(\mathcal{SCC}(z_{s,o}, \alpha))$$

  If we further consider tests $\mathcal{T}_\alpha$, $\alpha \in A^*$, obtained from $\mathcal{T}$ by making unsuccessful all the successful computations of $\mathcal{T}$ not compatible with α, we have that for each such test $prob(\mathcal{SC}^{\mathcal{T}_\alpha}(z_{s,o}))$ reduces to $prob(\mathcal{SCC}(z_{s,o}, \alpha))$. As a consequence, from $s_1 \sim_{\text{PTe-}\sqcup\sqcap} s_2$ we derive that for all $\alpha \in A^*$ it holds that $Res_{\max,\alpha}(s_1, o) \neq \emptyset$ iff $Res_{\max,\alpha}(s_2, o) \neq \emptyset$ and:

$$\bigsqcup_{\mathcal{Z}_1 \in Res_{\max}(s_1,o)} prob(\mathcal{SCC}(z_{s_1,o}, \alpha)) = \bigsqcup_{\mathcal{Z}_2 \in Res_{\max}(s_2,o)} prob(\mathcal{SCC}(z_{s_2,o}, \alpha))$$
$$\bigsqcap_{\mathcal{Z}_1 \in Res_{\max}(s_1,o)} prob(\mathcal{SCC}(z_{s_1,o}, \alpha)) = \bigsqcap_{\mathcal{Z}_2 \in Res_{\max}(s_2,o)} prob(\mathcal{SCC}(z_{s_2,o}, \alpha))$$

which means that $s_1 \sim_{\text{PTe-tbt},\sqcup\sqcap} s_2$. From this, it follows that $s_1 \sim^{\text{ct}}_{\text{PTe-tbt},\sqcup\sqcap} s_2$.

- Second, we show that $s_1 \sim^{\text{ct}}_{\text{PTe-tbt},\sqcup\sqcap} s_2 \implies s_1 \sim^{\text{ct}}_{\text{PTe-tbt}} s_2$. Suppose $s_1 \sim^{\text{ct}}_{\text{PTe-tbt},\sqcup\sqcap} s_2$ and consider an arbitrary trace $\alpha \in A^*$ for which there exists $\mathcal{Z}_1 \in \mathit{Res}^{\text{ct}}_{\max,\alpha}(s_1)$ such that $\mathit{prob}(\mathcal{SCC}(z_{s_1,o}, \alpha)) = p$. Since $s_1 \sim^{\text{ct}}_{\text{PTe-tbt},\sqcup\sqcap} s_2$, we have $\mathit{Res}^{\text{ct}}_{\max,\alpha}(s_2) \neq \emptyset$ and there exist $\mathcal{Z}'_2, \mathcal{Z}''_2 \in \mathit{Res}^{\text{ct}}_{\max,\alpha}(s_2)$ such that $\mathit{prob}(\mathcal{SCC}(z'_{s_2}, \alpha)) = p' \leq p$ and $\mathit{prob}(\mathcal{SCC}(z''_{s_2}, \alpha)) = p'' \geq p$.
  If $p' = p$ (resp. $p'' = p$), then $\mathcal{Z}_1$ is trivially matched by $\mathcal{Z}'_2$ (resp. $\mathcal{Z}''_2$) with respect to $\sim^{\text{ct}}_{\text{PTe-tbt}}$ when examining α.
  Assume that $p' < p < p''$ and consider the resolution $\mathcal{Z}_2 = x \cdot \mathcal{Z}'_2 + y \cdot \mathcal{Z}''_2$ of s_2 defined as follows for $x, y \in \mathbb{R}_{]0,1]}$ such that $x + y = 1$. Since $p' \neq p''$ and they both refer to the probability of performing a successful α-compatible computation from s_2, the two resolutions $\mathcal{Z}'_2$ and $\mathcal{Z}''_2$ of s_2 differ at least in one point in which the nondeterministic choice between two transitions labeled with the same action occurring in α has been resolved differently. We obtain $\mathcal{Z}_2$ from $\mathcal{Z}'_2$ and $\mathcal{Z}''_2$ by combining the two different transitions into a single one with coefficients x and y for their target distributions, respectively, in the first of those points. When examining α, if we take $x = \frac{p''-p}{p''-p'}$ and $y = \frac{p-p'}{p''-p'}$, then $\mathcal{Z}_1$ is matched by $\mathcal{Z}_2$ with respect to $\sim^{\text{ct}}_{\text{PTe-tbt}}$ because:
  $$\begin{aligned} \mathit{prob}(\mathcal{SCC}(z_{s_2,o}, \alpha)) &= \tfrac{p''-p}{p''-p'} \cdot \mathit{prob}(\mathcal{SCC}(z'_{s_2,o}, \alpha)) + \tfrac{p-p'}{p''-p'} \cdot \mathit{prob}(\mathcal{SCC}(z''_{s_2,o}, \alpha)) \\ &= \tfrac{p''-p}{p''-p'} \cdot p' + \tfrac{p-p'}{p''-p'} \cdot p'' \;=\; \tfrac{p' \cdot p'' - p \cdot p' + p \cdot p'' - p' \cdot p''}{p''-p'} \\ &= p \cdot \tfrac{p''-p'}{p''-p'} \;=\; p \;=\; \mathit{prob}(\mathcal{SCC}(z_{s_1,o}, \alpha)) \end{aligned}$$
  Due to the generality of $\alpha \in A^*$, it turns out that $s_1 \sim^{\text{ct}}_{\text{PTe-tbt}} s_2$.

Suppose now that $s_1 \sim_{\text{PTe-}\forall\exists} s_2$ and consider an arbitrary NPT $\mathcal{T} = (O, A, \longrightarrow)$ with initial state $o \in O$. Then, in particular, for all variants $\mathcal{T}_\alpha = (O, A, \longrightarrow_{\mathcal{T}_\alpha})$ of $\mathcal{T}$ in which all the successful computations of $\mathcal{T}$ not compatible with α are made unsuccessful, it holds that:

- For each $\mathcal{Z}_1 \in \mathit{Res}_{\max}(s_1, o)$ there exists $\mathcal{Z}_2 \in \mathit{Res}_{\max}(s_2, o)$ such that:
  $$\mathit{prob}(\mathcal{SC}^{\mathcal{T}_\alpha}(z_{s_1,o})) \;=\; \mathit{prob}(\mathcal{SC}^{\mathcal{T}_\alpha}(z_{s_2,o}))$$
- For each $\mathcal{Z}_2 \in \mathit{Res}_{\max}(s_2, o)$ there exists $\mathcal{Z}_1 \in \mathit{Res}_{\max}(s_1, o)$ such that:
  $$\mathit{prob}(\mathcal{SC}^{\mathcal{T}_\alpha}(z_{s_2,o})) \;=\; \mathit{prob}(\mathcal{SC}^{\mathcal{T}_\alpha}(z_{s_1,o}))$$

Since $\mathit{prob}(\mathcal{SC}^{\mathcal{T}_\alpha}(z_{s,o})) = \mathit{prob}(\mathcal{SCC}(z_{s,o}, \alpha))$ for all $s \in S$ due to the structure of $\mathcal{T}_\alpha$, we immediately derive that for all $\alpha \in A^*$ it holds that:

- For each $\mathcal{Z}_1 \in \mathit{Res}_{\max,\alpha}(s_1, o)$ there exists $\mathcal{Z}_2 \in \mathit{Res}_{\max,\alpha}(s_2, o)$ such that:
  $$\mathit{prob}(\mathcal{SCC}(z_{s_1,o}, \alpha)) \;=\; \mathit{prob}(\mathcal{SCC}(z_{s_2,o}, \alpha))$$
- For each $\mathcal{Z}_2 \in \mathit{Res}_{\max,\alpha}(s_2, o)$ there exists $\mathcal{Z}_1 \in \mathit{Res}_{\max,\alpha}(s_1, o)$ such that:
  $$\mathit{prob}(\mathcal{SCC}(z_{s_2,o}, \alpha)) \;=\; \mathit{prob}(\mathcal{SCC}(z_{s_1,o}, \alpha))$$

This means that $s_1 \sim_{\text{PTe-tbt}} s_2$. □

Apart from the use of $\mathit{prob}(\mathcal{SCC}(z_{s,o}, \alpha))$ values instead of $\mathit{prob}(\mathcal{SC}(z_{s,o}))$ values, another major difference between $\sim_{\text{PTe-tbt}}$ and $\sim_{\text{PTe-}\forall\exists}$ is the consideration of resolutions in $\mathit{Res}_{\max,\alpha}$ rather than in $\mathit{Res}_{\max}$. In other words, the considered maximal resolutions are those having at least one α-compatible computation that corresponds to a maximal α-compatible computation in the interaction system. The motivation behind this restriction is that it is not appropriate to match the 0 success probability of *maximal* α-compatible computations that are unsuccessful, with the 0 success probability of α-compatible computations that are *not maximal*, as may happen when considering $\mathit{Res}_{\max}$ instead of $\mathit{Res}_{\max,\alpha}$.

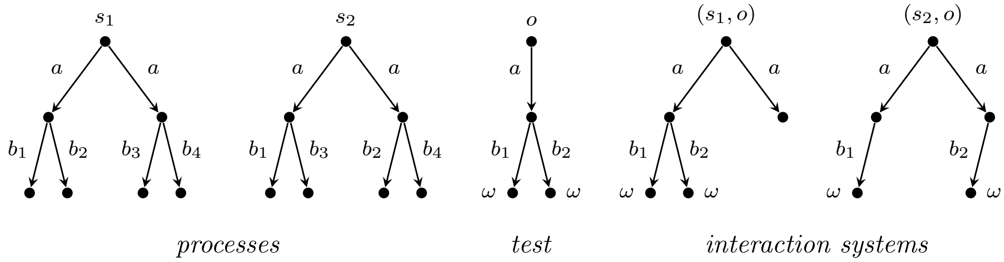

Figure 12: NPLTS models distinguished by $\sim_{\text{PTe-tbt}}$ thanks to the restriction to $Res_{\max,\alpha}$

Admitting all maximal resolutions would also cause $\sim_{\text{PTe-tbt}}$ not to be conservative with respect to $\sim_{\text{Te,fnd}}$ when restricting attention to fully nondeterministic tests. For example, if we consider the two fully nondeterministic NPLTS models in Fig. 12, it turns out that $s_1 \not\sim_{\text{Te,fnd}} s_2$ because of the fully nondeterministic test in the same figure. In fact, following the terminology of [11], the second process must pass that test, while the first one is not able to do so because the interaction system has a maximal computation labeled with a that does not reach success. In the setting of $\sim_{\text{PTe-tbt}}$, that computation in the first interaction system is not matched by any computation labeled with a in the second interaction system because of the restriction to $Res_{\max,a}$, thus correctly distinguishing the two processes. Notice that, under $Res_{\max}$, it would be matched by any of the two non-maximal computations labeled with a in the second interaction system.

We now investigate the inclusion and compatibility properties of $\sim_{\text{PTe-tbt}}/\sim^{\text{ct}}_{\text{PTe-tbt}}$. Similar to $\sim_{\text{PTe-}\forall\exists}$, they result in a testing semantics finer than trace semantics.

Theorem 5.3. *Let $(S, A, \longrightarrow)$ be an NPLTS and $s_1, s_2 \in S$. Then:*

$$s_1 \sim_{\text{PTe-tbt}} s_2 \implies s_1 \sim_{\text{PTr}} s_2$$
$$s_1 \sim^{\text{ct}}_{\text{PTe-tbt}} s_2 \implies s_1 \sim^{\text{ct}}_{\text{PTr}} s_2$$

Proof. If $s_1 \sim_{\text{PTe-tbt}} s_2$, then in particular for every NPT $\mathcal{T}_\alpha = (O, A, \longrightarrow_{\mathcal{T}_\alpha})$ with initial state $o \in O$ having a single maximal computation that is labeled with $\alpha \in A^*$ and reaches success, it holds that:

- For each $\mathcal{Z}_1 \in Res_{\max,\alpha}(s_1, o)$ there exists $\mathcal{Z}_2 \in Res_{\max,\alpha}(s_2, o)$ such that:
$$prob(\mathcal{SCC}(z_{s_1,o}, \alpha)) \;=\; prob(\mathcal{SCC}(z_{s_2,o}, \alpha))$$
- For each $\mathcal{Z}_2 \in Res_{\max,\alpha}(s_2, o)$ there exists $\mathcal{Z}_1 \in Res_{\max,\alpha}(s_1, o)$ such that:
$$prob(\mathcal{SCC}(z_{s_2,o}, \alpha)) \;=\; prob(\mathcal{SCC}(z_{s_1,o}, \alpha))$$

Since $prob(\mathcal{SCC}^{\mathcal{Z}}(z_{s,o}, \alpha)) = prob(\mathcal{CC}^{\mathcal{Z}'}(z_s, \alpha))$ for all $s \in S$ due to the structure of $\mathcal{T}_\alpha$ – where $\mathcal{Z} \in Res_{\max,\alpha}(s, o)$ and $\mathcal{Z}' \in Res(s)$ originates $\mathcal{Z}$ in the interaction with $\mathcal{T}_\alpha$ – we immediately derive that for all $\alpha \in A^*$ it holds that:

- For each $\mathcal{Z}_1 \in Res(s_1)$ there exists $\mathcal{Z}_2 \in Res(s_2)$ such that:
$$prob(\mathcal{CC}(z_{s_1}, \alpha)) \;=\; prob(\mathcal{CC}(z_{s_2}, \alpha))$$
- For each $\mathcal{Z}_2 \in Res(s_2)$ there exists $\mathcal{Z}_1 \in Res(s_1)$ such that:
$$prob(\mathcal{CC}(z_{s_2}, \alpha)) \;=\; prob(\mathcal{CC}(z_{s_1}, \alpha))$$

This means that $s_1 \sim_{\text{PTr}} s_2$.
The proof of $s_1 \sim^{\text{ct}}_{\text{PTe-tbt}} s_2 \implies s_1 \sim^{\text{ct}}_{\text{PTr}} s_2$ is analogous. $\square$

The inclusion of $\sim_{\text{PTe-tbt}}$ (resp. $\sim^{\text{ct}}_{\text{PTe-tbt}}$) in $\sim_{\text{PTr}}$ (resp. $\sim^{\text{ct}}_{\text{PTr}}$) is strict. For instance, the two NPLTS models in Fig. 7 are not trace-by-trace testing equivalent. In fact, the test in the same figure distinguishes s_1 from s_2 because – looking at the two interaction systems in Fig. 7 – each of the two maximal resolutions of (s_2, o) has a maximal computation labeled with a while the only maximal resolution of (s_1, o) has not.

Unlike $\sim_{\text{PTe-}\sqcup\sqcap}$ and $\sim_{\text{PTe-}\forall\exists}$, $\sim_{\text{PTe-tbt}}/\sim^{\text{ct}}_{\text{PTe-tbt}}$ result in a testing semantics that is *fully* (i.e., regardless of admitted tests) backward compatible with $\sim_{\text{Te,fnd}}$, $\sim_{\text{Te,fpr}}$, and $\sim_{\text{Te,rpr}}$. Concerning the two restricted classes of probabilistic processes, it is worth recalling that bisimulation equivalence and trace equivalence were defined uniformly for fully probabilistic processes [17, 23] and reactive probabilistic processes [26, 35]. In contrast, testing equivalence for fully probabilistic processes was defined in [8, 9] in a way that resembles $\sim_{\text{PTe-}\forall\exists}$, while for reactive probabilistic processes it was defined in [25] in a way similar to $\sim_{\text{PTe-}\sqcup\sqcap}$. Our compatibility results thus show that also testing equivalence could have been defined uniformly for both classes of probabilistic processes without internal nondeterminism, by resorting to the trace-by-trace approach that we have developed for NPLTS models.

Theorem 5.4. *Let $\mathcal{L} = (S, A, \longrightarrow)$ be an NPLTS and $s_1, s_2 \in S$.*

(1) *If $\mathcal{L}$ is fully nondeterministic, then:*
$$s_1 \sim_{\text{PTe-tbt}} s_2 \iff s_1 \sim^{\text{ct}}_{\text{PTe-tbt}} s_2 \iff s_1 \sim_{\text{Te,fnd}} s_2$$
(2) *If $\mathcal{L}$ is fully probabilistic, then:*
$$s_1 \sim_{\text{PTe-tbt}} s_2 \iff s_1 \sim^{\text{ct}}_{\text{PTe-tbt}} s_2 \iff s_1 \sim_{\text{Tr,fpr}} s_2$$
(3) *If $\mathcal{L}$ is reactive probabilistic, then:*
$$s_1 \sim_{\text{PTe-tbt}} s_2 \Longrightarrow s_1 \sim_{\text{Tr,rpr}} s_2$$
$$s_1 \sim^{\text{ct}}_{\text{PTe-tbt}} s_2 \Longrightarrow s_1 \sim_{\text{Tr,rpr}} s_2$$

Proof. We proceed as follows:

(1) Suppose that $\mathcal{L}$ is fully nondeterministic. We recall from [11] that $s_1 \sim_{\text{Te,fnd}} s_2$ means that for every *fully nondeterministic* NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$ it holds that:
   - There exists a successful computation from (s_1, o) iff there exists a successful computation from (s_2, o).
   - All maximal computations from (s_1, o) are successful iff all maximal computations from (s_2, o) are successful.

   In this setting, randomized schedulers are not important because, due to the absence of probabilistic choices, the model cannot contain submodels that arise from convex combinations of other submodels. Thus, we can concentrate on $\sim_{\text{PTe-tbt}}$. Suppose that $s_1 \sim_{\text{PTe-tbt}} s_2$. Then, in particular, for every *fully nondeterministic* NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$ and for all $\alpha \in A^*$ it holds that:
   - For each $\mathcal{Z}_1 \in Res_{\max,\alpha}(s_1, o)$ there exists $\mathcal{Z}_2 \in Res_{\max,\alpha}(s_2, o)$ such that:
   $$prob(\mathcal{SCC}(z_{s_1,o}, \alpha)) \;=\; prob(\mathcal{SCC}(z_{s_2,o}, \alpha))$$
   - For each $\mathcal{Z}_2 \in Res_{\max,\alpha}(s_2, o)$ there exists $\mathcal{Z}_1 \in Res_{\max,\alpha}(s_1, o)$ such that:
   $$prob(\mathcal{SCC}(z_{s_2,o}, \alpha)) \;=\; prob(\mathcal{SCC}(z_{s_1,o}, \alpha))$$

   Since the NPLTS under test and the considered tests are all fully nondeterministic, the resulting interaction systems are fully nondeterministic too, and hence their maximal resolutions coincide with their maximal computations and each of the probability values above is either 1 or 0. As a consequence, the previous relationships among maximal resolutions can be rephrased as follows:

- For each maximal α-compatible computation from (s_1, o) there exists a maximal α-compatible computation from (s_2, o) such that the two computations are both successful or both unsuccessful.
- For each maximal α-compatible computation from (s_2, o) there exists a maximal α-compatible computation from (s_1, o) such that the two computations are both successful or both unsuccessful.

From this, we immediately derive that:

- There exists a successful computation from (s_1, o) iff there exists a successful computation from (s_2, o).
- All maximal computations from (s_1, o) are successful iff all maximal computations from (s_2, o) are successful. In fact, assume that all maximal computations from, e.g., (s_1, o) are successful. Then at least one maximal computation from (s_2, o) is successful. Assume that (s_2, o) has at least two maximal computations and that one of them is not successful. Then at least one maximal computation from (s_1, o) would not be successful, thus contradicting the assumption that all maximal computations from (s_1, o) are successful. Therefore, whenever all maximal computations from (s_1, o) are successful, then all maximal computations from (s_2, o) are successful. Likewise, whenever all maximal computations from (s_2, o) are successful, then all maximal computations from (s_1, o) are successful.

This means that $s_1 \sim_{\mathrm{Te,fnd}} s_2$.

Suppose now that $s_1 \sim_{\mathrm{Te,fnd}} s_2$ and consider an *arbitrary* NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$, an arbitrary trace $\alpha \in A^*$ such that $\mathit{Res}_{\max,\alpha}(s_1, o) \neq \emptyset$, and an arbitrary resolution $\mathcal{Z}_1 \in \mathit{Res}_{\max,\alpha}(s_1, o)$.

Assume that $\mathit{Res}_{\max,\alpha}(s_2, o) = \emptyset$, i.e., assume that for all $\mathcal{Z}_2 \in \mathit{Res}_{\max}(s_2, o)$ it holds that $\mathcal{CC}_{\max}(z_{s_2,o}, \alpha) = \emptyset$. Let $\mathcal{T}_\alpha = (O, A, \longrightarrow_{\mathcal{T}_\alpha})$ be a *fully nondeterministic* NPT obtained from $\mathcal{T}$ in which (i) only the maximal α-compatible computations reach ω and (ii) each transition $o' \xrightarrow{a}_{\mathcal{T}} \mathcal{D}$ such that the set $O' = \{o'' \in O \mid \mathcal{D}(o'') > 0\}$ has cardinality greater than 1 is transformed into $|O'|$ transitions $o' \xrightarrow{a}_{\mathcal{T}_\alpha} \mathcal{D}_{o''}$, $o'' \in O'$, where $\mathcal{D}_{o''}(o'') = 1$ and $\mathcal{D}_{o''}(o''') = 0$ for all $o''' \in O \setminus \{o''\}$. Observing that $\mathcal{T}_\alpha$ yields the same α-compatible computations as $\mathcal{T}$ in the interaction systems, the test $\mathcal{T}_\alpha$ would violate $s_1 \sim_{\mathrm{Te,fnd}} s_2$ because at least one maximal computation from (s_1, o) is successful whilst there are no maximal computations from (s_2, o) that are successful. We have thus deduced that, whenever $s_1 \sim_{\mathrm{Te,fnd}} s_2$, then the existence of $\mathcal{Z}_1 \in \mathit{Res}_{\max,\alpha}(s_1, o)$ implies the existence of $\mathcal{Z}_2 \in \mathit{Res}_{\max,\alpha}(s_2, o)$.

Assume now that for all $\mathcal{Z}_2 \in \mathit{Res}_{\max,\alpha}(s_2, o)$ it holds that:

$$\mathit{prob}(\mathcal{SCC}(z_{s_1,o}, \alpha)) \neq \mathit{prob}(\mathcal{SCC}(z_{s_2,o}, \alpha))$$

Observing that $\mathcal{T}$ must have a successful α-compatible computation – otherwise it would hold that $\mathit{prob}(\mathcal{SCC}(z_{s_1,o}, \alpha)) = 0 = \mathit{prob}(\mathcal{SCC}(z_{s_2,o}, \alpha))$ for all $\mathcal{Z}_2 \in \mathit{Res}_{\max,\alpha}(s_2, o)$ – from $\mathcal{CC}_{\max}(z_{s_1,o}, \alpha) \neq \emptyset$ and $\mathcal{CC}_{\max}(z_{s_2,o}, \alpha) \neq \emptyset$ we derive that $\mathit{prob}(\mathcal{SCC}(z_{s_1,o}, \alpha)) > 0$ and $\mathit{prob}(\mathcal{SCC}(z_{s_2,o}, \alpha)) > 0$. Denoting by $\mathcal{Z}'_1$ the element of $\mathit{Res}_{\max}(s_1)$ that originates $\mathcal{Z}_1$, we would then have that for each $\mathcal{Z}'_2 \in \mathit{Res}_{\max}(s_2)$ originating $\mathcal{Z}_2$:

$$\begin{aligned} \mathit{prob}(\mathcal{CC}(z'_{s_1}, \alpha)) &= \mathit{prob}(\mathcal{SCC}(z_{s_1,o}, \alpha))/p \neq \\ &\neq \mathit{prob}(\mathcal{SCC}(z_{s_2,o}, \alpha))/p = \mathit{prob}(\mathcal{CC}(z'_{s_2}, \alpha)) \end{aligned}$$

where p is the probability of performing a successful α-compatible computation in the element $\mathcal{Z}$ of $\mathit{Res}_{\max}(o)$ that originates $\mathcal{Z}_1$. However, since the NPLTS under test is fully nondeterministic, $\mathcal{Z}'_1$ and $\mathcal{Z}'_2$ boil down to two α-compatible computations and it holds that:

$$prob(\mathcal{CC}(z'_{s_1}, \alpha)) \;=\; 1 \;=\; prob(\mathcal{CC}(z'_{s_2}, \alpha))$$

which contradicts what established before.

In conclusion, whenever $s_1 \sim_{\text{Te,fnd}} s_2$, then for each $\mathcal{Z}_1 \in Res_{\max,\alpha}(s_1, o)$ there exists $\mathcal{Z}_2 \in Res_{\max,\alpha}(s_2, o)$ such that:

$$prob(\mathcal{SCC}(z_{s_1,o}, \alpha)) \;=\; prob(\mathcal{SCC}(z_{s_2,o}, \alpha))$$

With a similar argument, we can prove that, whenever $s_1 \sim_{\text{Te,fnd}} s_2$, then for each $\mathcal{Z}_2 \in Res_{\max,\alpha}(s_2, o)$ there exists $\mathcal{Z}_1 \in Res_{\max,\alpha}(s_1, o)$ such that:

$$prob(\mathcal{SCC}(z_{s_2,o}, \alpha)) \;=\; prob(\mathcal{SCC}(z_{s_1,o}, \alpha))$$

This means that $s_1 \sim_{\text{PTe-tbt}} s_2$.

(2) Suppose that $\mathcal{L}$ is fully probabilistic. We recall from [9] that $s_1 \sim_{\text{Te,fpr}} s_2$ means that for every *fully probabilistic* NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$ it holds that:

$$prob(\mathcal{SC}(s_1, o)) \;=\; prob(\mathcal{SC}(s_2, o))$$

In this setting, schedulers are not important because there is no nondeterminism. Thus, we can concentrate on $\sim_{\text{PTe-tbt}}$. Suppose that $s_1 \sim_{\text{PTe-tbt}} s_2$. Then, in particular, for every *fully probabilistic* NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$ and for all $\alpha \in A^*$ it holds that:

- For each $\mathcal{Z}_1 \in Res_{\max,\alpha}(s_1, o)$ there exists $\mathcal{Z}_2 \in Res_{\max,\alpha}(s_2, o)$ such that:
$$prob(\mathcal{SCC}(z_{s_1,o}, \alpha)) \;=\; prob(\mathcal{SCC}(z_{s_2,o}, \alpha))$$
- For each $\mathcal{Z}_2 \in Res_{\max,\alpha}(s_2, o)$ there exists $\mathcal{Z}_1 \in Res_{\max,\alpha}(s_1, o)$ such that:
$$prob(\mathcal{SCC}(z_{s_2,o}, \alpha)) \;=\; prob(\mathcal{SCC}(z_{s_1,o}, \alpha))$$

Since the NPLTS under test and the considered tests are all fully probabilistic, the resulting interaction systems are fully probabilistic too, and hence each of them has a single maximal resolution that coincides with the interaction system itself. As a consequence, the previous relationships among maximal resolutions can be rephrased by saying that for all $\alpha \in A^*$:

$$prob(\mathcal{SCC}((s_1, o), \alpha)) \;=\; prob(\mathcal{SCC}((s_2, o), \alpha))$$

From this, we immediately derive that:

$$\begin{aligned} prob(\mathcal{SC}(s_1, o)) \;&=\; \sum_{\alpha \in A^*} prob(\mathcal{SCC}((s_1, o), \alpha)) \;=\; \\ &=\; \sum_{\alpha \in A^*} prob(\mathcal{SCC}((s_2, o), \alpha)) \;=\; prob(\mathcal{SC}(s_2, o)) \end{aligned}$$

which means that $s_1 \sim_{\text{Te,fpr}} s_2$.

Suppose now that $s_1 \sim_{\text{Te,fpr}} s_2$ and consider an *arbitrary* NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$, an arbitrary trace $\alpha \in A^*$ such that $Res_{\max,\alpha}(s_1, o) \neq \emptyset$, and an arbitrary resolution $\mathcal{Z}_1 \in Res_{\max,\alpha}(s_1, o)$.

Assume that $Res_{\max,\alpha}(s_2, o) = \emptyset$, i.e., assume that for all $\mathcal{Z}_2 \in Res_{\max}(s_2, o)$ it holds that $\mathcal{CC}_{\max}(z_{s_2,o}, \alpha) = \emptyset$. Let $\mathcal{T}_\alpha = (O, A, \longrightarrow_{\mathcal{T}_\alpha})$ be a *fully probabilistic* NPT obtained from $\mathcal{T}$ in which (i) only the maximal α-compatible computations reach ω, (ii) each state $o' \in O$ having at most one outgoing transition $o' \xrightarrow{a}_{\mathcal{T}} \mathcal{D}$ retains all of its transitions, and (iii) any other state in O retains among its transitions only one of those that are instrumental to preserve the original α-compatible computations of $\mathcal{T}$. Observing that $\mathcal{T}_\alpha$ yields at least one of the α-compatible computations of $\mathcal{T}$ in the interaction systems, the test $\mathcal{T}_\alpha$ would violate $s_1 \sim_{\text{Te,fpr}} s_2$ because at least one maximal computation from (s_1, o) is successful whilst there are no maximal computations from (s_2, o) that are successful. We have thus deduced that, whenever $s_1 \sim_{\text{Te,fpr}} s_2$, then the existence of $\mathcal{Z}_1 \in Res_{\max,\alpha}(s_1, o)$ implies the existence of $\mathcal{Z}_2 \in Res_{\max,\alpha}(s_2, o)$.

Assume now that for all $\mathcal{Z}_2 \in Res_{\max,\alpha}(s_2, o)$ it holds that:

$$prob(\mathcal{SCC}(z_{s_1,o},\alpha)) \neq prob(\mathcal{SCC}(z_{s_2,o},\alpha))$$

Observing that $\mathcal{T}$ must have a successful α-compatible computation – otherwise it would hold that $prob(\mathcal{SCC}(z_{s_1,o},\alpha)) = 0 = prob(\mathcal{SCC}(z_{s_2,o},\alpha))$ for all $\mathcal{Z}_2 \in Res_{\max,\alpha}(s_2, o)$ – from $\mathcal{CC}_{\max}(z_{s_1,o},\alpha) \neq \emptyset$ and $\mathcal{CC}_{\max}(z_{s_2,o},\alpha) \neq \emptyset$ we derive that $prob(\mathcal{SCC}(z_{s_1,o},\alpha)) > 0$ and $prob(\mathcal{SCC}(z_{s_2,o},\alpha)) > 0$. Denoting by $\mathcal{Z}'_1$ the element of $Res_{\max}(s_1)$ that originates $\mathcal{Z}_1$, we would then have that for each $\mathcal{Z}'_2 \in Res_{\max}(s_2)$ originating $\mathcal{Z}_2$:

$$\begin{aligned} prob(\mathcal{CC}(z'_{s_1},\alpha)) \;=\; prob(\mathcal{SCC}(z_{s_1,o},\alpha))/p \neq \\ \neq prob(\mathcal{SCC}(z_{s_2,o},\alpha))/p \;=\; prob(\mathcal{CC}(z'_{s_2},\alpha)) \end{aligned}$$

where p is the probability of performing a successful α-compatible computation in the element $\mathcal{Z}$ of $Res_{\max}(o)$ that originates $\mathcal{Z}_1$. However, since the NPLTS under test is fully probabilistic, it holds that:

$$\begin{aligned} prob(\mathcal{CC}(z'_{s_1},\alpha)) \;&=\; prob(\mathcal{CC}(s_1,\alpha)) \\ prob(\mathcal{CC}(z'_{s_2},\alpha)) \;&=\; prob(\mathcal{CC}(s_2,\alpha)) \end{aligned}$$

where:

$$prob(\mathcal{CC}(s_1,\alpha)) \;=\; prob(\mathcal{CC}(s_2,\alpha))$$

because otherwise $s_1 \sim_{\mathrm{Te,fpr}} s_2$ would be violated by a test having a single maximal computation that is labeled with α and reaches ω. Thus:

$$prob(\mathcal{CC}(z'_{s_1},\alpha)) \;=\; prob(\mathcal{CC}(z'_{s_2},\alpha))$$

which contradicts what established before.

In conclusion, whenever $s_1 \sim_{\mathrm{Te,fpr}} s_2$, then for each $\mathcal{Z}_1 \in Res_{\max,\alpha}(s_1, o)$ there exists $\mathcal{Z}_2 \in Res_{\max,\alpha}(s_2, o)$ such that:

$$prob(\mathcal{SCC}(z_{s_1,o},\alpha)) \;=\; prob(\mathcal{SCC}(z_{s_2,o},\alpha))$$

With a similar argument, we can prove that, whenever $s_1 \sim_{\mathrm{Te,fpr}} s_2$, then for each $\mathcal{Z}_2 \in Res_{\max,\alpha}(s_2, o)$ there exists $\mathcal{Z}_1 \in Res_{\max,\alpha}(s_1, o)$ such that:

$$prob(\mathcal{SCC}(z_{s_2,o},\alpha)) \;=\; prob(\mathcal{SCC}(z_{s_1,o},\alpha))$$

This means that $s_1 \sim_{\mathrm{PTe\text{-}tbt}} s_2$.

(3) Suppose that $\mathcal{L}$ is reactive probabilistic. Taking inspiration from [25], $s_1 \sim_{\mathrm{Te,rpr}} s_2$ means that for every *reactive probabilistic* NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$ it holds that:

$$\begin{aligned} \bigsqcup_{\alpha \in Tr_{\max}(s_1,o)} prob(\mathcal{SCC}((s_1,o),\alpha)) \;&=\; \bigsqcup_{\alpha \in Tr_{\max}(s_2,o)} prob(\mathcal{SCC}((s_2,o),\alpha)) \\ \bigsqcap_{\alpha \in Tr_{\max}(s_1,o)} prob(\mathcal{SCC}((s_1,o),\alpha)) \;&=\; \bigsqcap_{\alpha \in Tr_{\max}(s_2,o)} prob(\mathcal{SCC}((s_2,o),\alpha)) \end{aligned}$$

Given $s \in S$, the set $Tr_{\max}(s, o)$ contains all the traces labeling the maximal computations from (s, o), while success probabilities $prob(\mathcal{SCC}((s,o),\alpha))$ are viewed as being conditional on selecting the maximal resolution of (s, o) that contains all the α-compatible computations from (s, o) (this resolution is unique because interaction systems are reactive probabilistic).

Suppose that $s_1 \sim_{\mathrm{PTe\text{-}tbt}} s_2$. Then, in particular, for every *reactive probabilistic* NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$ and for all $\alpha \in A^*$ it holds that:

- For each $\mathcal{Z}_1 \in Res_{\max,\alpha}(s_1, o)$ there exists $\mathcal{Z}_2 \in Res_{\max,\alpha}(s_2, o)$ such that:
$$prob(\mathcal{SCC}(z_{s_1,o},\alpha)) \;=\; prob(\mathcal{SCC}(z_{s_2,o},\alpha))$$
- For each $\mathcal{Z}_2 \in Res_{\max,\alpha}(s_2, o)$ there exists $\mathcal{Z}_1 \in Res_{\max,\alpha}(s_1, o)$ such that:
$$prob(\mathcal{SCC}(z_{s_2,o},\alpha)) \;=\; prob(\mathcal{SCC}(z_{s_1,o},\alpha))$$

Since the NPLTS under test and the considered tests are all reactive probabilistic, the resulting interaction systems are reactive probabilistic too, and hence in each of them there is a unique maximal resolution that collects all the computations compatible with a given maximal trace. As a consequence, from the previous relationships among maximal resolutions we derive that for all $\alpha \in A^*$:

$$prob(\mathcal{SCC}((s_1, o), \alpha)) \; = \; prob(\mathcal{SCC}((s_2, o), \alpha))$$

From this, we immediately derive that:

$$\bigsqcup_{\alpha \in \mathit{Tr}_{\max}(s_1,o)} prob(\mathcal{SCC}((s_1, o), \alpha)) \; = \; \bigsqcup_{\alpha \in \mathit{Tr}_{\max}(s_2,o)} prob(\mathcal{SCC}((s_2, o), \alpha))$$

$$\bigsqcap_{\alpha \in \mathit{Tr}_{\max}(s_1,o)} prob(\mathcal{SCC}((s_1, o), \alpha)) \; = \; \bigsqcap_{\alpha \in \mathit{Tr}_{\max}(s_2,o)} prob(\mathcal{SCC}((s_2, o), \alpha))$$

which means that $s_1 \sim_{\text{Te,rpr}} s_2$.

The proof that $s_1 \sim^{\text{ct}}_{\text{PTe-tbt}} s_2$ implies $s_1 \sim_{\text{Te,rpr}} s_2$ is similar. □

In [12], it was shown that $\sim_{\text{PTe-}\sqcup\sqcap}$ is a congruence with respect to parallel composition. To conclude, we prove that also the trace-by-trace approach results in a compositional testing semantics.

Theorem 5.5. *Let $\mathcal{L}_k = (S_k, A, \longrightarrow_k)$ be an NPLTS for $k = 0, 1, 2$ and consider $\mathcal{L}_1 \,\|_{\mathcal{A}}\, \mathcal{L}_0$ and $\mathcal{L}_2 \,\|_{\mathcal{A}}\, \mathcal{L}_0$ for $\mathcal{A} \subseteq A$. Let $s_k \in S_k$ for $k = 0, 1, 2$. Then:*

$$s_1 \sim_{\text{PTe-tbt}} s_2 \implies (s_1, s_0) \sim_{\text{PTe-tbt}} (s_2, s_0)$$

$$s_1 \sim^{\text{ct}}_{\text{PTe-tbt}} s_2 \implies (s_1, s_0) \sim^{\text{ct}}_{\text{PTe-tbt}} (s_2, s_0)$$

Proof. Given an arbitrary NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$, first of all we observe that $\mathcal{L}_0 \,\|\, \mathcal{T}$ is still an NPT, with initial state $(s_0, o) \in S_0 \times O$.

If $s_1 \sim_{\text{PTe-tbt}} s_2$, then in particular for all $\alpha \in A^*$ it holds that:

- For each $\mathcal{Z}_1 \in \mathit{Res}_{\max,\alpha}(s_1, (s_0, o))$ there exists $\mathcal{Z}_2 \in \mathit{Res}_{\max,\alpha}(s_2, (s_0, o))$ such that:
$$prob(\mathcal{SCC}(z_{s_1,(s_0,o)}, \alpha)) \; = \; prob(\mathcal{SCC}(z_{s_2,(s_0,o)}, \alpha))$$
- For each $\mathcal{Z}_2 \in \mathit{Res}_{\max,\alpha}(s_2, (s_0, o))$ there exists $\mathcal{Z}_1 \in \mathit{Res}_{\max,\alpha}(s_1, (s_0, o))$ such that:
$$prob(\mathcal{SCC}(z_{s_2,(s_0,o)}, \alpha)) \; = \; prob(\mathcal{SCC}(z_{s_1,(s_0,o)}, \alpha))$$

For $h = 1, 2$, we note that $(s_h, (s_0, o))$ is a configuration of $\mathcal{L}_h \,\|\, (\mathcal{L}_0 \,\|\, \mathcal{T})$ while $((s_h, s_0), o)$ is a configuration of $(\mathcal{L}_h \,\|_{\mathcal{A}}\, \mathcal{L}_0) \,\|\, \mathcal{T}$, hence $\mathit{Res}_{\max,\alpha}(s_h, (s_0, o)) \subseteq \mathit{Res}_{\max,\alpha}((s_h, s_0), o)$ because $\mathcal{L}_h \,\|\, (\mathcal{L}_0 \,\|\, \mathcal{T})$ is fully synchronous. There are three cases.

If $\mathcal{A} = A$, then $(\mathcal{L}_h \,\|_{\mathcal{A}}\, \mathcal{L}_0) \,\|\, \mathcal{T} = (\mathcal{L}_h \,\|\, \mathcal{L}_0) \,\|\, \mathcal{T}$ and we can exploit associativity of $\|$ to establish that $\mathit{Res}_{\max,\alpha}(s_h, (s_0, o)) = \mathit{Res}_{\max,\alpha}((s_h, s_0), o)$ for $h = 1, 2$.

If $\mathcal{A} \subset A$ and $\mathcal{L}_1$ and $\mathcal{L}_2$ have no transitions labeled with actions not in $\mathcal{A}$, then for $h = 1, 2$ it holds that all transitions of $\mathcal{L}_h$ must synchronize with transitions of $\mathcal{T}$ both in $\mathcal{L}_h \,\|\, (\mathcal{L}_0 \,\|\, \mathcal{T})$ and in $(\mathcal{L}_h \,\|_{\mathcal{A}}\, \mathcal{L}_0) \,\|\, \mathcal{T}$, hence possible resolutions in $\mathit{Res}_{\max,\alpha}((s_h, s_0), o)$ that do not belong to $\mathit{Res}_{\max,\alpha}(s_h, (s_0, o))$ are due to transitions of $\mathcal{L}_0$ not labeled with actions in $\mathcal{A}$ that synchronize with transitions of $\mathcal{T}$.

If $\mathcal{A} \subset A$ and $\mathcal{L}_1$ and $\mathcal{L}_2$ have transitions labeled with actions not in $\mathcal{A}$, then these transitions (which originate resolutions in $\mathit{Res}_{\max,\alpha}((s_h,s_0), o)$ that do not belong to $\mathit{Res}_{\max,\alpha}(s_h, (s_0,o))$ for $h = 1, 2$) must occur in corresponding points of $\mathcal{L}_1$ and $\mathcal{L}_2$ (otherwise we could find a test that distinguishes s_1 from s_2 with respect to $\sim_{\text{PTe-tbt}}$) and must synchronize with transitions of $\mathcal{T}$ in order for them to emerge in the interaction systems.

In each of the three cases, for all $\alpha \in A^*$ it holds that:

- For each $\mathcal{Z}_1 \in \mathit{Res}_{\max,\alpha}((s_1, s_0), o)$ there exists $\mathcal{Z}_2 \in \mathit{Res}_{\max,\alpha}((s_2, s_0), o)$ such that:
$$prob(\mathcal{SCC}(z_{(s_1,s_0),o}, \alpha)) \; = \; prob(\mathcal{SCC}(z_{(s_2,s_0),o}, \alpha))$$
- For each $\mathcal{Z}_2 \in \mathit{Res}_{\max,\alpha}((s_2, s_0), o)$ there exists $\mathcal{Z}_1 \in \mathit{Res}_{\max,\alpha}((s_1, s_0), o)$ such that:
$$prob(\mathcal{SCC}(z_{(s_2,s_0),o}, \alpha)) \; = \; prob(\mathcal{SCC}(z_{(s_1,s_0),o}, \alpha))$$

This means that $(s_1, s_0) \sim_{\text{PTe-tbt}} (s_2, s_0)$ because $\mathcal{T}$ is an arbitrary NPT.
The proof of compositionality for $\sim^{\text{ct}}_{\text{PTe-tbt}}$ is analogous. $\square$

6. Placing Trace and Testing Equivalences in a Spectrum

In this section, we investigate the relationships between the various equivalences that we have recalled from the literature ($\sim_{\text{PTr,dis}}$ and $\sim_{\text{PTe-}\sqcup\sqcap}$) or introduced for the first time ($\sim_{\text{PTr}}$, $\sim_{\text{PTe-}\forall\exists}$, and $\sim_{\text{PTe-tbt}}$) together with their variants based on randomized schedulers. Some inclusion, coincidence, and incomparability results have already been established in Thms. 3.3, 4.4, 4.6, 4.7, 5.2, and 5.3.

We start by providing a surprising characterization of the finest relation considered so far, i.e., $\sim_{\text{PTe-}\forall\exists}$, that will be useful later on to establish a connection with failure semantics. The characterization is expressed in terms of a variant of $\sim_{\text{PTe-tbt}}$, denoted by $\sim_{\text{PTe-tbt,dis}}$, that is inspired by $\sim_{\text{PTr,dis}}$ and hence considers successful trace distributions.

Definition 6.1. Let $(S, A, \longrightarrow)$ be an NPLTS. We say that $s_1, s_2 \in S$ are *probabilistic trace-by-trace-distribution testing equivalent*, written $s_1 \sim_{\text{PTe-tbt,dis}} s_2$, iff for every NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$ it holds that:

- For each $\mathcal{Z}_1 \in \mathit{Res}_{\max}(s_1, o)$ there exists $\mathcal{Z}_2 \in \mathit{Res}_{\max}(s_2, o)$ such that <u>for all $\alpha \in A^*$</u> it holds that $\mathcal{CC}_{\max}(z_{s_1,o}, \alpha) \neq \emptyset$ implies $\mathcal{CC}_{\max}(z_{s_2,o}, \alpha) \neq \emptyset$ and:
$$\mathit{prob}(\mathcal{SCC}(z_{s_1,o}, \alpha)) \;=\; \mathit{prob}(\mathcal{SCC}(z_{s_2,o}, \alpha))$$
- For each $\mathcal{Z}_2 \in \mathit{Res}_{\max}(s_2, o)$ there exists $\mathcal{Z}_1 \in \mathit{Res}_{\max}(s_1, o)$ such that <u>for all $\alpha \in A^*$</u> it holds that $\mathcal{CC}_{\max}(z_{s_2,o}, \alpha) \neq \emptyset$ implies $\mathcal{CC}_{\max}(z_{s_1,o}, \alpha) \neq \emptyset$ and:
$$\mathit{prob}(\mathcal{SCC}(z_{s_2,o}, \alpha)) \;=\; \mathit{prob}(\mathcal{SCC}(z_{s_1,o}, \alpha))$$

We denote by $\sim^{\text{ct}}_{\text{PTe-tbt,dis}}$ the coarser variant based on randomized schedulers. ■

Theorem 6.2. *Let $(S, A, \longrightarrow)$ be an NPLTS and $s_1, s_2 \in S$. Then:*
$$s_1 \sim_{\text{PTe-}\forall\exists} s_2 \iff s_1 \sim_{\text{PTe-tbt,dis}} s_2$$
$$s_1 \sim^{\text{ct}}_{\text{PTe-}\forall\exists} s_2 \iff s_1 \sim^{\text{ct}}_{\text{PTe-tbt,dis}} s_2$$

Proof. Let us prove the contrapositive of $s_1 \sim_{\text{PTe-}\forall\exists} s_2 \implies s_1 \sim_{\text{PTe-tbt,dis}} s_2$. Thus, suppose that $s_1 \not\sim_{\text{PTe-tbt,dis}} s_2$. This means that there exist an NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$ and, say, a resolution $\mathcal{Z}_1 \in \mathit{Res}_{\max}(s_1, o)$ such that for each $\mathcal{Z}_2 \in \mathit{Res}_{\max}(s_2, o)$ there exists $\alpha_2 \in A^*$ such that $\mathcal{CC}_{\max}(z_{s_1,o}, \alpha_2) \neq \emptyset$ and (i) $\mathcal{CC}_{\max}(z_{s_2,o}, \alpha_2) = \emptyset$ or (ii) $\mathit{prob}(\mathcal{SCC}(z_{s_1,o}, \alpha_2)) \neq \mathit{prob}(\mathcal{SCC}(z_{s_2,o}, \alpha_2))$. We show that from this fact it follows that $s_1 \not\sim_{\text{PTe-}\forall\exists} s_2$ by proceeding by induction on the number n of traces labeling the successful computations from o (note that n is finite – because $\mathcal{T}$ is finite – and greater than 0 – otherwise $\mathcal{T}$ cannot distinguish s_1 from s_2 with respect to $\sim_{\text{PTe-tbt,dis}}$):

- Let $n = 1$ and denote by α the only trace labeling the successful computations from o. Then $\mathcal{CC}_{\max}(z_{s_1,o}, \alpha) \neq \emptyset$ and (i) $\mathcal{CC}_{\max}(z_{s_2,o}, \alpha) = \emptyset$ in which case:
$$\mathit{prob}(\mathcal{SC}(z_{s_1,o})) \;>\; 0 \;=\; \mathit{prob}(\mathcal{SC}(z_{s_2,o}))$$
or (ii) it holds that:
$$\begin{aligned}\mathit{prob}(\mathcal{SC}(z_{s_1,o})) \;=\; \mathit{prob}(\mathcal{SCC}(z_{s_1,o}, \alpha)) \;\neq\; \\ \neq\; \mathit{prob}(\mathcal{SCC}(z_{s_2,o}, \alpha)) \;=\; \mathit{prob}(\mathcal{SC}(z_{s_2,o}))\end{aligned}$$
As a consequence, in both cases $s_1 \not\sim_{\text{PTe-}\forall\exists} s_2$.
- Let $n \in \mathbb{N}_{>1}$ and suppose that the result holds for all $m = 1, \ldots, n - 1$. Given a trace α labeling some of the successful computations from o, we denote by $\mathcal{T}_{\downarrow\alpha}$ the NPT obtained

from $\mathcal{T}$ by transforming into a normal terminal state every success state reached by a maximal α-compatible computation, and by $\mathcal{T}_{\uparrow\alpha}$ the NPT obtained from $\mathcal{T}$ by transforming into a normal terminal state every success state reached by a maximal computation not compatible with α. Since $\mathcal{T}$ distinguishes s_1 from s_2 with respect to $\sim_{\text{PTe-tbt,dis}}$, $\mathcal{T}_{\downarrow\alpha}$ and $\mathcal{T}_{\uparrow\alpha}$ have the same structure as $\mathcal{T}$, and α labels some of the successful computations of $\mathcal{T}$, either $\mathcal{T}_{\downarrow\alpha}$ or $\mathcal{T}_{\uparrow\alpha}$ still distinguishes s_1 from s_2 with respect to $\sim_{\text{PTe-tbt,dis}}$. Since $\mathcal{T}_{\downarrow\alpha}$ has $n-1$ traces labeling its successful computations and $\mathcal{T}_{\uparrow\alpha}$ has a single trace labeling its successful computations, by the induction hypothesis it follows that $s_1 \not\sim_{\text{PTe-}\forall\exists} s_2$.

Suppose now that $s_1 \sim_{\text{PTe-tbt,dis}} s_2$ and consider an arbitrary NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$. Since for all $s \in S$ and $\mathcal{Z} \in \mathit{Res}_{\max}(s, o)$ it holds that:

$$\mathit{prob}(\mathcal{SC}(z_{s,o})) \; = \sum_{\alpha \in A^* \text{ s.t. } \mathcal{CC}_{\max}(z_{s,o},\alpha) \neq \emptyset} \mathit{prob}(\mathcal{SCC}(z_{s,o}, \alpha))$$

from $s_1 \sim_{\text{PTe-tbt,dis}} s_2$ it follows that:

- For each $\mathcal{Z}_1 \in \mathit{Res}_{\max}(s_1, o)$ there exists $\mathcal{Z}_2 \in \mathit{Res}_{\max}(s_2, o)$ such that:
$$\begin{aligned}\mathit{prob}(\mathcal{SC}(z_{s_1,o})) \; &= \sum_{\alpha \in A^* \text{ s.t. } \mathcal{CC}_{\max}(z_{s_1,o},\alpha) \neq \emptyset} \mathit{prob}(\mathcal{SCC}(z_{s_1,o}, \alpha)) \; = \\ &= \sum_{\alpha \in A^* \text{ s.t. } \mathcal{CC}_{\max}(z_{s_2,o},\alpha) \neq \emptyset} \mathit{prob}(\mathcal{SCC}(z_{s_2,o}, \alpha)) \; = \; \mathit{prob}(\mathcal{SC}(z_{s_2,o}))\end{aligned}$$
- For each $\mathcal{Z}_2 \in \mathit{Res}_{\max}(s_2, o)$ there exists $\mathcal{Z}_1 \in \mathit{Res}_{\max}(s_1, o)$ such that:
$$\begin{aligned}\mathit{prob}(\mathcal{SC}(z_{s_2,o})) \; &= \sum_{\alpha \in A^* \text{ s.t. } \mathcal{CC}_{\max}(z_{s_2,o},\alpha) \neq \emptyset} \mathit{prob}(\mathcal{SCC}(z_{s_2,o}, \alpha)) \; = \\ &= \sum_{\alpha \in A^* \text{ s.t. } \mathcal{CC}_{\max}(z_{s_1,o},\alpha) \neq \emptyset} \mathit{prob}(\mathcal{SCC}(z_{s_1,o}, \alpha)) \; = \; \mathit{prob}(\mathcal{SC}(z_{s_1,o}))\end{aligned}$$

This means that $s_1 \sim_{\text{PTe-}\forall\exists} s_2$.
The fact that $\sim^{\text{ct}}_{\text{PTe-}\forall\exists}$ and $\sim^{\text{ct}}_{\text{PTe-tbt,dis}}$ coincide immediately follows. □

We know from [10] that for fully nondeterministic processes there is a strong connection between the testing semantics of [11] and the failure semantics of [5]. Thus, for a more complete comparison of the various trace and testing equivalences, we also present failure semantics for NPLTS models. In particular, we consider two variants $\sim_{\text{PF,dis}}/\sim^{\text{ct}}_{\text{PF,dis}}$ of the probabilistic failure-distribution equivalence defined in [33] on the basis of the pattern of $\sim^{\text{ct}}_{\text{PTr,dis}}$ [32], and we introduce two variants $\sim_{\text{PF}}/\sim^{\text{ct}}_{\text{PF}}$ of a novel probabilistic failure equivalence by taking inspiration from the pattern of $\sim_{\text{PTr}}$. We shall see that $\sim_{\text{PTe-}\forall\exists}$ (i.e., $\sim_{\text{PTe-tbt,dis}}$) is strictly finer than $\sim_{\text{PF,dis}}/\sim^{\text{ct}}_{\text{PF,dis}}$, while $\sim_{\text{PTe-tbt}}$ and $\sim^{\text{ct}}_{\text{PTe-tbt}}$ are strictly coarser than $\sim_{\text{PF}}$ and $\sim^{\text{ct}}_{\text{PF}}$, respectively.

In the following, we call *failure pair* an element $\varphi \in A^* \times 2^A$ formed by a trace α and a failure set F. Given a state s of an NPLTS $\mathcal{L}$, a resolution $\mathcal{Z}$ of s, and a computation $c \in \mathcal{C}_{\text{fin}}(z_s)$, we say that c is compatible with φ iff $c \in \mathcal{CC}(z_s, \alpha)$ and the state in $\mathcal{L}$ corresponding to the last state reached by c has no outgoing transitions in $\mathcal{L}$ labeled with an action in F. We denote by $\mathcal{FCC}(z_s, \varphi)$ the set of φ-compatible computations from z_s.

Definition 6.3. Let $(S, A, \longrightarrow)$ be an NPLTS. We say that $s_1, s_2 \in S$ are *probabilistic failure-distribution equivalent*, written $s_1 \sim_{\text{PF,dis}} s_2$, iff:

- For each $\mathcal{Z}_1 \in \mathit{Res}(s_1)$ there exists $\mathcal{Z}_2 \in \mathit{Res}(s_2)$ such that $\underline{\text{for all } \varphi \in A^* \times 2^A}$:
$$\mathit{prob}(\mathcal{FCC}(z_{s_1}, \varphi)) \; = \; \mathit{prob}(\mathcal{FCC}(z_{s_2}, \varphi))$$
- For each $\mathcal{Z}_2 \in \mathit{Res}(s_2)$ there exists $\mathcal{Z}_1 \in \mathit{Res}(s_1)$ such that $\underline{\text{for all } \varphi \in A^* \times 2^A}$:
$$\mathit{prob}(\mathcal{FCC}(z_{s_2}, \varphi)) \; = \; \mathit{prob}(\mathcal{FCC}(z_{s_1}, \varphi))$$

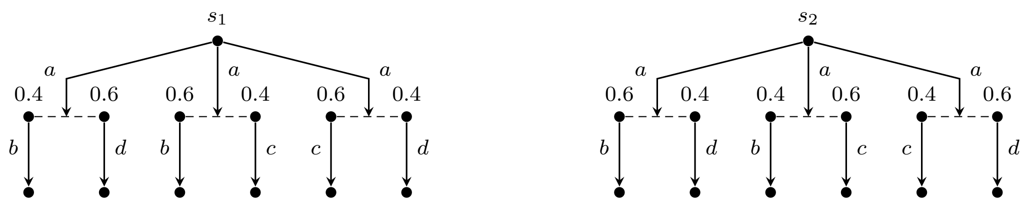

Figure 13: NPLTS models distinguished by $\sim_{\rm PF,dis}/\sim_{\rm PF,dis}^{\rm ct}$ and identified by $\sim_{\rm PF}/\sim_{\rm PF}^{\rm ct}$

We denote by $\sim_{\rm PF,dis}^{\rm ct}$ the coarser variant based on randomized schedulers. ■

Definition 6.4. Let $(S, A, \longrightarrow)$ be an NPLTS. We say that $s_1, s_2 \in S$ are *probabilistic failure equivalent*, written $s_1 \sim_{\rm PF} s_2$, iff $\underline{\text{for all } \varphi \in A^* \times 2^A}$ it holds that:

- For each $\mathcal{Z}_1 \in Res(s_1)$ there exists $\mathcal{Z}_2 \in Res(s_2)$ such that:
$$prob(\mathcal{FCC}(z_{s_1}, \varphi)) \;=\; prob(\mathcal{FCC}(z_{s_2}, \varphi))$$
- For each $\mathcal{Z}_2 \in Res(s_2)$ there exists $\mathcal{Z}_1 \in Res(s_1)$ such that:
$$prob(\mathcal{FCC}(z_{s_2}, \varphi)) \;=\; prob(\mathcal{FCC}(z_{s_1}, \varphi))$$

We denote by $\sim_{\rm PF}^{\rm ct}$ the coarser variant based on randomized schedulers. ■

Theorem 6.5. *Let $(S, A, \longrightarrow)$ be an NPLTS and $s_1, s_2 \in S$. Then:*
$$s_1 \sim_{\rm PF,dis} s_2 \implies s_1 \sim_{\rm PF} s_2$$
$$s_1 \sim_{\rm PF,dis}^{\rm ct} s_2 \implies s_1 \sim_{\rm PF}^{\rm ct} s_2$$

Proof. If $s_1 \sim_{\rm PF,dis} s_2$ (resp. $s_1 \sim_{\rm PF,dis}^{\rm ct} s_2$), then $s_1 \sim_{\rm PF} s_2$ (resp. $s_1 \sim_{\rm PF}^{\rm ct} s_2$) follows by taking the same fully matching resolutions considered for $\sim_{\rm PF,dis}$ (resp. $\sim_{\rm PF,dis}^{\rm ct}$). □

The inclusion of $\sim_{\rm PF,dis}$ (resp. $\sim_{\rm PF,dis}^{\rm ct}$) in $\sim_{\rm PF}$ (resp. $\sim_{\rm PF}^{\rm ct}$) is strict, because the initial states of the two NPLTS models in Fig. 13 are equated by the latter equivalence and told apart by the former. Moreover, Figs. 13 and 5 together show that $\sim_{\rm PF}$ and $\sim_{\rm PF,dis}^{\rm ct}$ are incomparable with each other.

Theorem 6.6. *Let $(S, A, \longrightarrow)$ be an NPLTS and $s_1, s_2 \in S$. Then:*
$$s_1 \sim_{\rm PTe\text{-}tbt,dis} s_2 \implies s_1 \sim_{\rm PF,dis} s_2$$

Proof. Firstly, we prove that $s_1 \sim_{\rm PTe\text{-}tbt,dis} s_2 \implies s_1 \sim_{\rm PRTr,dis} s_2$ where $\sim_{\rm PRTr,dis}$ is defined as follows. We call *ready trace* an element $\rho \in (A \times 2^A)^*$ given by a sequence of $n \in \mathbb{N}$ pairs of the form (a_i, R_i). Given $s \in S$, $\mathcal{Z} \in Res(s)$, and $c \in \mathcal{C}_{\rm fin}(z_s)$, we say that c is compatible with ρ iff $c \in \mathcal{CC}(z_s, a_1 \dots a_n)$ and, denoting by z_i the state reached by c after the i-th step for all $i = 1, \dots, n$, the set of actions labeling the transitions in $\mathcal{L}$ departing from the state in $\mathcal{L}$ corresponding to z_i is precisely R_i. We denote by $\mathcal{RTCC}(z_s, \rho)$ the set of ρ-compatible computations from z_s. We say that s_1 and s_2 are *probabilistic ready-trace-distribution equivalent*, written $s_1 \sim_{\rm PRTr,dis} s_2$, iff for each $\mathcal{Z}_1 \in Res(s_1)$ there exists $\mathcal{Z}_2 \in Res(s_2)$ such that for all $\rho \in (A \times 2^A)^*$:
$$prob(\mathcal{RTCC}(z_{s_1}, \rho)) \;=\; prob(\mathcal{RTCC}(z_{s_2}, \rho))$$
and symmetrically for each $\mathcal{Z}_2 \in Res(s_2)$.
We show that $s_1 \sim_{\rm PTe\text{-}tbt,dis} s_2$ implies $s_1 \sim_{\rm PRTr,dis} s_2$ by building a test that permits to reason about all ready traces at once for each resolution of s_1 and s_2. We start by deriving a new NPLTS $(S_{\rm r}, A_{\rm r}, \longrightarrow_{\rm r})$ that is isomorphic to the given one up to transition labels

and terminal states. A transition $s \xrightarrow{a} \mathcal{D}$ becomes $s_{\mathrm{r}} \xrightarrow{a \triangleleft R}_{\mathrm{r}} \mathcal{D}_{\mathrm{r}}$ where $R \subseteq A$ is the set of actions labeling the outgoing transitions of s and $\mathcal{D}_{\mathrm{r}}(s_{\mathrm{r}}) = \mathcal{D}(s)$ for all $s \in S$. If s is a terminal state, i.e., it has no outgoing transitions, then we add a transition $s_{\mathrm{r}} \xrightarrow{\circ \triangleleft \emptyset}_{\mathrm{r}} \delta_{s_{\mathrm{r}}}$ where $\delta_{s_{\mathrm{r}}}(s_{\mathrm{r}}) = 1$ and $\delta_{s_{\mathrm{r}}}(s'_{\mathrm{r}}) = 0$ for all $s' \in S \setminus \{s\}$. Transition relabeling preserves $\sim_{\mathrm{PTe\text{-}tbt,dis}}$, i.e., $s_1 \sim_{\mathrm{PTe\text{-}tbt,dis}} s_2$ implies $s_{1,\mathrm{r}} \sim_{\mathrm{PTe\text{-}tbt,dis}} s_{2,\mathrm{r}}$, because $\sim_{\mathrm{PTe\text{-}tbt,dis}}$ is able to distinguish a state that has a single α-compatible computation reaching a state with a nondeterministic branching formed by a b-transition and a c-transition, from a state that has two α-compatible computations such that one of them reaches a state with only one outgoing transition labeled with b and the other one reaches a state with only one outgoing transition labeled with c (e.g., use a test that has a single α-compatible computation whose last step leads to a distribution whose support contains only a state with only one outgoing transition labeled with b that reaches success and a state with only one outgoing transition labeled with c that reaches success).

For each $\alpha_{\mathrm{r}} \in (A_{\mathrm{r}})^*$ and $R \subseteq A$, we build an NPT $\mathcal{T}_{\alpha_{\mathrm{r}},R} = (O_{\alpha_{\mathrm{r}},R}, A_{\mathrm{r}}, \longrightarrow_{\alpha_{\mathrm{r}},R})$ having a single α_{r}-compatible computation that goes from the initial state $o_{\alpha_{\mathrm{r}},R}$ to a state having a single transition to ω labeled with (i) $\circ \triangleleft \emptyset$ if $R = \emptyset$ or (ii) $\_ \triangleleft R$ if $R \neq \emptyset$. Since we compare individual states (like s_1 and s_2) rather than state distributions, the distinguishing power of $\sim_{\mathrm{PTe\text{-}tbt,dis}}$ does not change if we additionally consider tests starting with a single τ-transition that can initially evolve autonomously in any interaction system. We thus build a further NPT $\mathcal{T} = (O, A_{\mathrm{r}}, \longrightarrow_{\mathcal{T}})$ that has an initial τ-transition and then behaves as one of the tests $\mathcal{T}_{\alpha_{\mathrm{r}},R}$, i.e., its initial τ-transition goes from the initial state o to a state distribution whose support is the set $\{o_{\alpha_{\mathrm{r}},R} \mid \alpha_{\mathrm{r}} \in (A_{\mathrm{r}})^* \wedge R \subseteq A\}$, with the probability $p_{\alpha_{\mathrm{r}},R}$ associated with $o_{\alpha_{\mathrm{r}},R}$ being taken from the distribution whose values are of the form $1/2^i$, $i \in \mathbb{N}_{>0}$. Note that $\mathcal{T}$ is not finite state, but this affects only the initial step, whose only purpose is to internally select a specific ready trace.

After this step, $\mathcal{T}$ interacts with the process under test. Let $\rho \in (A \times 2^A)^*$ be a ready trace of the form $(a_1, R_1) \dots (a_n, R_n)$, where $n \in \mathbb{N}$. Given $s \in S$, consider the trace $\alpha_{\rho,\mathrm{r}} \in (A_{\mathrm{r}})^*$ of length $n+1$ in which the first element is $a_1 \triangleleft R$, with $R \subseteq A$ being the set of actions labeling the outgoing transitions of s, the subsequent elements are of the form $a_i \triangleleft R_{i-1}$ for $i = 2, \dots, n$, and the last element is (i) $\circ \triangleleft \emptyset$ if $R_n = \emptyset$ or (ii) $\_ \triangleleft R_n$ if $R_n \neq \emptyset$. Then for all $\mathcal{Z} \in \mathit{Res}(s)$ it holds that:

$$\mathit{prob}(\mathcal{RTCC}(z_s, \rho)) \; = \; 0$$

if there is no $a_1 \dots a_n$-compatible computation from z_s, otherwise:

$$\mathit{prob}(\mathcal{RTCC}(z_s, \rho)) \; = \; \mathit{prob}(\mathcal{SCC}(z_{s_{\mathrm{r}},o}, \alpha_{\rho,\mathrm{r}})) / p_{\alpha'_{\rho,\mathrm{r}},R_n}$$

where $\alpha'_{\rho,\mathrm{r}}$ is $\alpha_{\rho,\mathrm{r}}$ without its last element.

Suppose that $s_1 \sim_{\mathrm{PTe\text{-}tbt,dis}} s_2$, which implies that s_1 and s_2 have the same set R of actions labeling their outgoing transitions and $s_{1,\mathrm{r}} \sim_{\mathrm{PTe\text{-}tbt,dis}} s_{2,\mathrm{r}}$. Then:

- For each $\mathcal{Z}_1 \in \mathit{Res}(s_1)$ there exists $\mathcal{Z}_2 \in \mathit{Res}(s_2)$ such that for all ready traces $\rho = (a_1, R_1) \dots (a_n, R_n) \in (A \times 2^A)^*$ either:
$$\mathit{prob}(\mathcal{RTCC}(z_{s_1}, \rho)) \; = \; 0 \; = \; \mathit{prob}(\mathcal{RTCC}(z_{s_2}, \rho))$$
or:
$$\begin{aligned} \mathit{prob}(\mathcal{RTCC}(z_{s_1}, \rho)) \; &= \; \mathit{prob}(\mathcal{SCC}(z_{s_{1,\mathrm{r}},o}, \alpha_{\rho,\mathrm{r}})) / p_{\alpha'_{\rho,\mathrm{r}},R_n} \; = \\ &= \; \mathit{prob}(\mathcal{SCC}(z_{s_{2,\mathrm{r}},o}, \alpha_{\rho,\mathrm{r}})) / p_{\alpha'_{\rho,\mathrm{r}},R_n} \; = \; \mathit{prob}(\mathcal{RTCC}(z_{s_2}, \rho)) \end{aligned}$$
- Symmetrically for each $\mathcal{Z}_2 \in \mathit{Res}(s_2)$.

This means that $s_1 \sim_{\mathrm{PRTr,dis}} s_2$.

Secondly, we prove that $s_1 \sim_{\mathrm{PRTr,dis}} s_2 \implies s_1 \sim_{\mathrm{PFTr,dis}} s_2$ where $\sim_{\mathrm{PFTr,dis}}$ is defined as follows. We call *failure trace* an element $\phi \in (A \times 2^A)^*$ given by a sequence of $n \in \mathbb{N}$ pairs of the form (a_i, F_i). Given $s \in S$, $\mathcal{Z} \in \mathit{Res}(s)$, and $c \in \mathcal{C}_{\mathrm{fin}}(z_s)$, we say that c is compatible with ϕ iff $c \in \mathcal{CC}(z_s, a_1 \dots a_n)$ and, denoting by z_i the state reached by c after the i-th step for all $i = 1, \dots, n$, the state in $\mathcal{L}$ corresponding to z_i has no outgoing transitions in $\mathcal{L}$ labeled with an action in F_i. We denote by $\mathcal{FTCC}(z_s, \phi)$ the set of ϕ-compatible computations from z_s. We say that s_1 and s_2 are *probabilistic failure-trace-distribution equivalent*, written $s_1 \sim_{\mathrm{PFTr,dis}} s_2$, iff for each $\mathcal{Z}_1 \in \mathit{Res}(s_1)$ there exists $\mathcal{Z}_2 \in \mathit{Res}(s_2)$ such that for all $\phi \in (A \times 2^A)^*$:

$$\mathit{prob}(\mathcal{FTCC}(z_{s_1}, \phi)) \; = \; \mathit{prob}(\mathcal{FTCC}(z_{s_2}, \phi))$$

and symmetrically for each $\mathcal{Z}_2 \in \mathit{Res}(s_2)$.

Suppose that $s_1 \sim_{\mathrm{PRTr,dis}} s_2$. Since for all $s \in S$, $\mathcal{Z} \in \mathit{Res}(s)$, $n \in \mathbb{N}$, $\alpha = a_1 \dots a_n \in A^*$, and $F_1, \dots, F_n \in 2^A$ it holds that:

$$\mathit{prob}(\mathcal{FTCC}(z_s, (a_1, F_1) \dots (a_n, F_n))) \; = \; \sum_{R'_1, \dots, R'_n \in 2^A \text{ s.t. } R'_i \cap F_i = \emptyset \text{ for all } i = 1, \dots, n} \mathit{prob}(\mathcal{RTCC}(z_s, (a_1, R'_1) \dots (a_n, R'_n)))$$

we immediately derive that:

- For each $\mathcal{Z}_1 \in \mathit{Res}(s_1)$ there exists $\mathcal{Z}_2 \in \mathit{Res}(s_2)$ such that for all failure traces $(a_1, F_1) \dots (a_n, F_n) \in (A \times 2^A)^*$:

$$\begin{aligned} & \mathit{prob}(\mathcal{FTCC}(z_{s_1}, (a_1, F_1) \dots (a_n, F_n))) \; = \\ & \quad = \; \textstyle\sum_{R'_1, \dots, R'_n \in 2^A \text{ s.t. } R'_i \cap F_i = \emptyset \text{ for all } i = 1, \dots, n} \mathit{prob}(\mathcal{RTCC}(z_{s_1}, (a_1, R'_1) \dots (a_n, R'_n))) \\ & \quad = \; \textstyle\sum_{R'_1, \dots, R'_n \in 2^A \text{ s.t. } R'_i \cap F_i = \emptyset \text{ for all } i = 1, \dots, n} \mathit{prob}(\mathcal{RTCC}(z_{s_2}, (a_1, R'_1) \dots (a_n, R'_n))) \\ & \quad = \; \mathit{prob}(\mathcal{FTCC}(z_{s_2}, (a_1, F_1) \dots (a_n, F_n))) \end{aligned}$$

- Symmetrically for each $\mathcal{Z}_2 \in \mathit{Res}(s_2)$.

This means that $s_1 \sim_{\mathrm{PFTr,dis}} s_2$.

Thirdly, we prove that $s_1 \sim_{\mathrm{PFTr,dis}} s_2 \Longrightarrow s_1 \sim_{\mathrm{PF,dis}} s_2$. Suppose that $s_1 \sim_{\mathrm{PFTr,dis}} s_2$. Since for all $s \in S$, $\mathcal{Z} \in \mathit{Res}(s)$, $n \in \mathbb{N}$, $\alpha = a_1 \dots a_n \in A^*$, and $F \in 2^A$ it holds that:

$$\mathit{prob}(\mathcal{FCC}(z_s, (\alpha, F))) \; = \; \mathit{prob}(\mathcal{FTCC}(z_s, (a_1, \emptyset) \dots (a_{n-1}, \emptyset)(a_n, F)))$$

we immediately derive that:

- For each $\mathcal{Z}_1 \in \mathit{Res}(s_1)$ there exists $\mathcal{Z}_2 \in \mathit{Res}(s_2)$ such that for all failure pairs $(a_1 \dots a_n, F) \in A^* \times 2^A$:

$$\begin{aligned} \mathit{prob}(\mathcal{FCC}(z_{s_1}, (a_1 \dots a_n, F))) \; & = \; \mathit{prob}(\mathcal{FTCC}(z_{s_1}, (a_1, \emptyset) \dots (a_{n-1}, \emptyset)(a_n, F))) \\ & = \; \mathit{prob}(\mathcal{FTCC}(z_{s_2}, (a_1, \emptyset) \dots (a_{n-1}, \emptyset)(a_n, F))) \\ & = \; \mathit{prob}(\mathcal{FCC}(z_{s_2}, (a_1 \dots a_n, F))) \end{aligned}$$

- Symmetrically for each $\mathcal{Z}_2 \in \mathit{Res}(s_2)$.

This means that $s_1 \sim_{\mathrm{PF,dis}} s_2$. □

The inclusion of $\sim_{\mathrm{PTe\text{-}tbt,dis}}$ in $\sim_{\mathrm{PF,dis}}$ is strict, because for the two NPLTS models in Fig. 8 it holds that $s_1 \sim_{\mathrm{PF,dis}} s_2$ while $s_1 \not\sim_{\mathrm{PTe\text{-}tbt,dis}} s_2$ as witnessed by the test in the same figure (see the maximal resolutions of the interaction systems in Fig. 9).

Theorem 6.7. *Let* $(S, A, \longrightarrow)$ *be an NPLTS and* $s_1, s_2 \in S$. *Then:*

$$s_1 \sim_{\mathrm{PF,dis}} s_2 \Longrightarrow s_1 \sim_{\mathrm{PTr,dis}} s_2$$
$$s_1 \sim^{\mathrm{ct}}_{\mathrm{PF,dis}} s_2 \Longrightarrow s_1 \sim^{\mathrm{ct}}_{\mathrm{PTr,dis}} s_2$$

Proof. Suppose that $s_1 \sim_{\mathrm{PF,dis}} s_2$. Then $s_1 \sim_{\mathrm{PTr,dis}} s_2$ because for all $s \in S$, $\mathcal{Z} \in \mathit{Res}(s)$, and $\alpha \in A^*$ it holds that:

$$\mathit{prob}(\mathcal{CC}(z_s, \alpha)) \; = \; \mathit{prob}(\mathcal{FCC}(z_s, (\alpha, \emptyset)))$$

and hence:

- For each $\mathcal{Z}_1 \in \mathit{Res}(s_1)$ there exists $\mathcal{Z}_2 \in \mathit{Res}(s_2)$ such that for all $\alpha \in A^*$:
$$\begin{aligned} \mathit{prob}(\mathcal{CC}(z_{s_1}, \alpha)) \;=\; & \mathit{prob}(\mathcal{FCC}(z_{s_1}, (\alpha, \emptyset))) \;= \\ = \; & \mathit{prob}(\mathcal{FCC}(z_{s_2}, (\alpha, \emptyset))) \;=\; \mathit{prob}(\mathcal{CC}(z_{s_2}, \alpha)) \end{aligned}$$
- Symmetrically for each $\mathcal{Z}_2 \in \mathit{Res}(s_2)$.

The proof that $s_1 \sim^{\rm ct}_{\rm PF,dis} s_2$ implies $s_1 \sim^{\rm ct}_{\rm PTr,dis} s_2$ is similar. □

The inclusion of $\sim_{\rm PF,dis}$ (resp. $\sim^{\rm ct}_{\rm PF,dis}$) in $\sim_{\rm PTr,dis}$ (resp. $\sim^{\rm ct}_{\rm PTr,dis}$) is strict, because the initial states of the two NPLTS models in Fig. 7 are equated by the latter equivalence and told apart by the former.

Theorem 6.8. *Let $(S, A, \longrightarrow)$ be an NPLTS and $s_1, s_2 \in S$. Then:*
$$\begin{aligned} s_1 \sim_{\rm PF} s_2 &\implies s_1 \sim_{\rm PTe\text{-}tbt} s_2 \\ s_1 \sim^{\rm ct}_{\rm PF} s_2 &\implies s_1 \sim^{\rm ct}_{\rm PTe\text{-}tbt} s_2 \end{aligned}$$

Proof. Let us prove the contrapositive of the first result, i.e., $s_1 \not\sim_{\rm PTe\text{-}tbt} s_2 \implies s_1 \not\sim_{\rm PF} s_2$. Thus, suppose that $s_1 \not\sim_{\rm PTe\text{-}tbt} s_2$. This means that there exist an NPT $\mathcal{T} = (O, A, \longrightarrow_{\mathcal{T}})$ with initial state $o \in O$, a trace $\alpha \in A^*$, and, say, a resolution $\mathcal{Z}_1 \in \mathit{Res}_{\max,\alpha}(s_1, o)$ such that $\mathit{Res}_{\max,\alpha}(s_2, o) = \emptyset$ or for all $\mathcal{Z}_2 \in \mathit{Res}_{\max,\alpha}(s_2, o)$ it holds that:
$$\mathit{prob}(\mathcal{SCC}(z_{s_1,o}, \alpha)) \;\neq\; \mathit{prob}(\mathcal{SCC}(z_{s_2,o}, \alpha))$$
Observing that $\mathit{Res}_{\max,\alpha}(s_1, o) \neq \emptyset$, in the case that $\mathit{Res}_{\max,\alpha}(s_2, o) = \emptyset$ either s_2 cannot perform α at all – let $\varphi = (\alpha, \emptyset)$ – or, after performing α, the states reached by s_2 can always synchronize with the states reached by o on a set F of actions whereas the states reached by s_1 cannot – let $\varphi = (\alpha, F)$. The failure pair φ shows that $s_1 \not\sim_{\rm PF} s_2$ in this case because, denoting by $\mathcal{Z}'_1$ the element of $\mathit{Res}(s_1)$ that originates $\mathcal{Z}_1$, we have that for all $\mathcal{Z}'_2 \in \mathit{Res}(s_2)$:
$$\mathit{prob}(\mathcal{FCC}(z'_{s_1}, \varphi)) \;>\; 0 \;=\; \mathit{prob}(\mathcal{FCC}(z'_{s_2}, \varphi))$$
In the case that $\mathit{Res}_{\max,\alpha}(s_2, o) \neq \emptyset$, the failure pair $\varphi = (\alpha, \emptyset)$ shows that $s_1 \not\sim_{\rm PF} s_2$. In fact, without loss of generality we can assume that the only α-compatible computations in $\mathcal{T}$ are the ones exercised by $\mathcal{Z}_1$ – note that they must belong to the same element $\mathcal{Z}$ of $\mathit{Res}(o)$ – as the only effect of this assumption is that of possibly reducing the number of resolutions in $\mathit{Res}_{\max,\alpha}(s_2, o)$. At least one of these computations must be successful – and hence maximal – in $\mathcal{T}$ because otherwise the success probabilities of the considered resolutions would all be equal to 0. Denoting by $\mathcal{Z}'_1$ the element of $\mathit{Res}(s_1)$ that originates $\mathcal{Z}_1$, we then have that for all $\mathcal{Z}'_2 \in \mathit{Res}(s_2)$ originating some $\mathcal{Z}_2 \in \mathit{Res}_{\max,\alpha}(s_2, o)$:
$$\begin{aligned} \mathit{prob}(\mathcal{FCC}(z'_{s_1}, \varphi)) \;=\; & \mathit{prob}(\mathcal{SCC}(z_{s_1,o}, \alpha))/p \;\neq \\ \neq \; & \mathit{prob}(\mathcal{SCC}(z_{s_2,o}, \alpha))/p \;=\; \mathit{prob}(\mathcal{FCC}(z'_{s_2}, \varphi)) \end{aligned}$$
where p is the probability of performing the α-compatible computations in the only element $\mathcal{Z}$ of $\mathit{Res}(o)$ that originates $\mathcal{Z}_1$ and all the resolutions $\mathcal{Z}_2$.

The proof that $s_1 \sim^{\rm ct}_{\rm PF} s_2$ implies $s_1 \sim^{\rm ct}_{\rm PTe\text{-}tbt} s_2$ is similar. □

The inclusion of $\sim_{\rm PF}$ (resp. $\sim^{\rm ct}_{\rm PF}$) in $\sim_{\rm PTe\text{-}tbt}$ (resp. $\sim^{\rm ct}_{\rm PTe\text{-}tbt}$) is strict, because the initial states of the two NPLTS models in Fig. 4 are equated by the latter equivalence and told apart by the former. For instance, the rightmost maximal resolution of s_1 has probability 1 of performing a computation compatible with the failure pair $(a, \{b_1, b_2\})$, whilst each of the two maximal resolutions of s_2 has probability 0.5.

The relationships among the various probabilistic testing, failure, and trace equivalences for NPLTS models are summarized in Fig. 14. Arrows represent the more-discriminating-than partial order, equivalences close to each other coincide, and incomparability is denoted

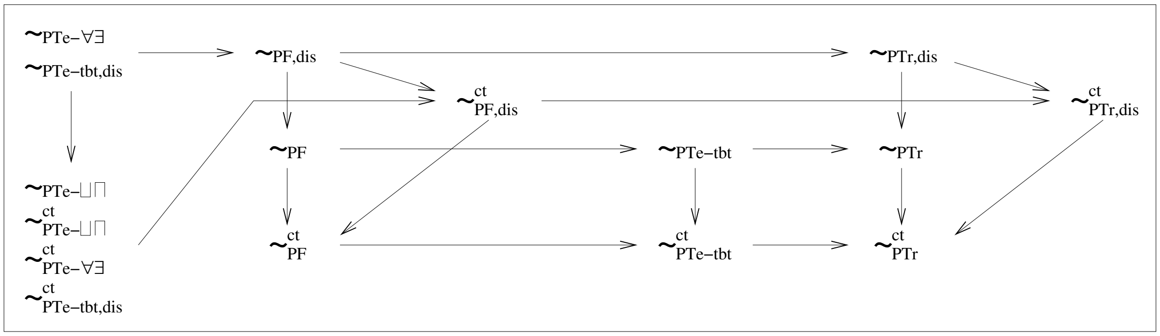

Figure 14: The spectrum of testing, failure, and trace equivalences for NPLTS models

by the absence of (chains of) arrows. The various relationships have been established in this paper, except for the arrow from $\sim_{\text{PTe-}\sqcup\sqcap}$ to $\sim^{\text{ct}}_{\text{PF,dis}}$ that is due to [33].

We observe that $\sim_{\text{PTe-}\sqcup\sqcap}$ is incomparable not only with $\sim_{\text{PTe-tbt}}$ as established right before Thm. 5.2, but also with $\sim_{\text{PF,dis}}$, $\sim_{\text{PF}}$, $\sim_{\text{PTr,dis}}$, and $\sim_{\text{PTr}}$. In fact, in Fig. 5 it holds that $s_1 \sim_{\text{PTe-}\sqcup\sqcap} s_2$ while $s_1 \not\sim_{\text{PF,dis}} s_2$, $s_1 \not\sim_{\text{PF}} s_2$, $s_1 \not\sim_{\text{PTr,dis}} s_2$, and $s_1 \not\sim_{\text{PTr}} s_2$. On the other hand, in Fig. 8 it holds that $s_1 \not\sim_{\text{PTe-}\sqcup\sqcap} s_2$ while $s_1 \sim_{\text{PF,dis}} s_2$, $s_1 \sim_{\text{PF}} s_2$, $s_1 \sim_{\text{PTr,dis}} s_2$, and $s_1 \sim_{\text{PTr}} s_2$.

Likewise, $\sim^{\text{ct}}_{\text{PF,dis}}$ is incomparable not only with $\sim_{\text{PF}}$ as established right after Thm. 6.5, but also with $\sim_{\text{PTr,dis}}$, $\sim_{\text{PTe-tbt}}$, and $\sim_{\text{PTr}}$. Indeed, in Fig. 5 it holds that $s_1 \sim^{\text{ct}}_{\text{PF,dis}} s_2$ while $s_1 \not\sim_{\text{PTr,dis}} s_2$, $s_1 \not\sim_{\text{PTe-tbt}} s_2$, and $s_1 \not\sim_{\text{PTr}} s_2$. In contrast, in Fig. 13 it holds that $s_1 \not\sim^{\text{ct}}_{\text{PF,dis}} s_2$ while $s_1 \sim_{\text{PTr,dis}} s_2$, $s_1 \sim_{\text{PTe-tbt}} s_2$, and $s_1 \sim_{\text{PTr}} s_2$. Moreover, $\sim^{\text{ct}}_{\text{PF}}$ is incomparable with $\sim_{\text{PTr,dis}}$ and $\sim^{\text{ct}}_{\text{PTr,dis}}$. In fact, in Fig. 13 it holds that $s_1 \sim^{\text{ct}}_{\text{PF}} s_2$ while $s_1 \not\sim_{\text{PTr,dis}} s_2$ and $s_1 \not\sim^{\text{ct}}_{\text{PTr,dis}} s_2$. On the other hand, in Fig. 7 it holds that $s_1 \not\sim^{\text{ct}}_{\text{PF}} s_2$ while $s_1 \sim_{\text{PTr,dis}} s_2$ and $s_1 \sim^{\text{ct}}_{\text{PTr,dis}} s_2$. Additionally, $\sim^{\text{ct}}_{\text{PF}}$ is incomparable with $\sim_{\text{PTr}}$ because in Fig. 5 we have that $s_1 \sim^{\text{ct}}_{\text{PF}} s_2$ and $s_1 \not\sim_{\text{PTr}} s_2$, whereas in Fig. 7 we have that $s_1 \not\sim^{\text{ct}}_{\text{PF}} s_2$ and $s_1 \sim_{\text{PTr}} s_2$. Furthermore, $\sim^{\text{ct}}_{\text{PF}}$ is incomparable also with $\sim_{\text{PTe-tbt}}$ because in Fig. 5 we have that $s_1 \sim^{\text{ct}}_{\text{PF}} s_2$ and $s_1 \not\sim_{\text{PTe-tbt}} s_2$, whilst in Fig. 4 we have that $s_1 \not\sim^{\text{ct}}_{\text{PF}} s_2$ and $s_1 \sim_{\text{PTe-tbt}} s_2$.

Analogously, $\sim^{\text{ct}}_{\text{PTr,dis}}$ is incomparable not only with $\sim_{\text{PTr}}$ as established right after Thm. 3.3, but also with $\sim_{\text{PF}}$, $\sim_{\text{PTe-tbt}}$, and $\sim^{\text{ct}}_{\text{PTe-tbt}}$. It holds that $s_1 \sim^{\text{ct}}_{\text{PTr,dis}} s_2$ and $s_1 \not\sim_{\text{PF}} s_2$, $s_1 \not\sim_{\text{PTe-tbt}} s_2$, and $s_1 \not\sim^{\text{ct}}_{\text{PTe-tbt}} s_2$ in Fig. 7, while $s_1 \not\sim^{\text{ct}}_{\text{PTr,dis}} s_2$ and $s_1 \sim_{\text{PF}} s_2$, $s_1 \sim_{\text{PTe-tbt}} s_2$, and $s_1 \sim^{\text{ct}}_{\text{PTe-tbt}} s_2$ in Fig. 13. The same two figures show that also $\sim_{\text{PTr,dis}}$ is incomparable with $\sim_{\text{PF}}$, $\sim_{\text{PTe-tbt}}$, and $\sim^{\text{ct}}_{\text{PTe-tbt}}$. Finally, we have that $\sim_{\text{PTr}}$ is incomparable with $\sim^{\text{ct}}_{\text{PTe-tbt}}$ because in Fig. 7 it holds that $s_1 \sim_{\text{PTr}} s_2$ and $s_1 \not\sim^{\text{ct}}_{\text{PTe-tbt}} s_2$, whereas in Fig. 5 it holds that $s_1 \not\sim_{\text{PTr}} s_2$ and $s_1 \sim^{\text{ct}}_{\text{PTe-tbt}} s_2$.

We conclude by recalling another probabilistic testing equivalence that has been recently proposed in [16], where a probabilistic model significantly different from ours is considered. Unfortunately, the differences prevent us from placing that equivalence in the spectrum we have just presented. However, that testing equivalence shares with our $\sim_{\text{PTe-tbt}}$ motivations and intuitions concerning the power of schedulers and the estimation of success probabilities that call for further comments.

The model considered in [16] has three types of transitions: action transitions, internal transitions, and probabilistic transitions. Since each state can have only one type of outgoing

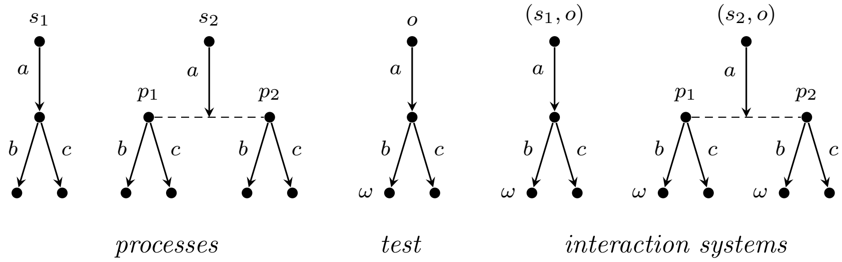

Figure 15: NPLTS models equated by [16] and distinguished by $\sim_{\text{PTe-tbt}}$

transitions, also states are divided into three classes: action states, nondeterministic states, and probabilistic states. Action states cannot have two identically labeled action transitions, so this model can be viewed as a variant of reactive probabilistic processes in which states of different classes can alternate along a computation. Notice that our NPLTS model is non-alternating, because there is a single class of states and probabilistic choices are somehow embedded within each single transition.

In order to make the proposed testing theory insensitive to the exact moment in which internal choices occur, in [16] internal transitions are decorated with so-called internal labels. Similar to action states, nondeterministic states cannot have two identically labeled internal transitions. Moreover, given two nondeterministic states, either they share the same set of internal labels decorating their outgoing transitions, or the sets of internal labels of their outgoing transitions are disjoint. Internal labels are meant to provide precisely the information that schedulers should use to resolve internal choices, so that internal choices relying on the same information are resolved in the same way. For example, continuing the discussion done in the last two paragraphs of Sect. 4, with the approach of [16] the two internal choices between the two b-transitions in the interaction system with initial configuration (s_1, o) of Figs. 8 and 10 would be identically tagged, say with b_l and b_r based on the orientation of the arrows. As a consequence, the only allowed maximal resolutions of that interaction system among the four shown in Figs. 9 and 11 would be the first one (choice of b_l) and the fourth one (choice of b_r), thus excluding success probabilities 1 and 0.

An important technical point made in [16] is that, in the presence of cycles of transitions within the model, the same internal choice may occur several times along a computation. This is not due to the copying capability that arises when composing in parallel a process and a test, which – as we have recalled above – is dealt with by labeling in the same way the internal transitions departing from all the copies of the cloned state and by forcing schedulers to perform consistent choices in all the copies (we will refer to the resulting fully probabilistic models as *consistent resolutions*). Replications of the same internal choice at different *unfolding depths* of a cycle are independent of each other and are thus given additional labels that keep them distinct from depth to depth. Notice that, in contrast, our approach based on $\sim_{\text{PTe-tbt}}$ is not invasive at all, as it does not require any label massaging on the model to restrict the power of schedulers.

Two processes are equated by the testing equivalence proposed in [16] iff, for each test, every consistent resolution at unfolding depth m of a suitably labeled version of the first interaction system that reaches success with probability p, is matched by a consistent resolution at the same unfolding depth of a suitably labeled version of the second interaction system that reaches success with the same probability. This equivalence cannot be directly

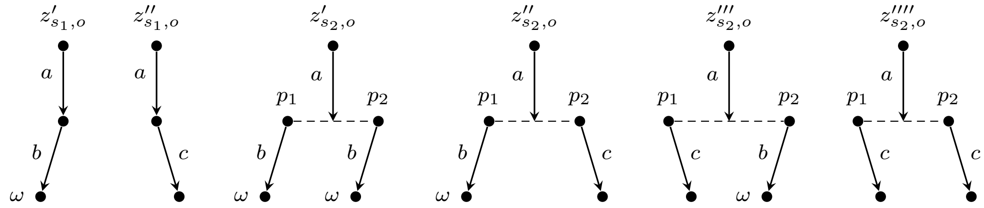

Figure 16: Maximal resolutions of the two interaction systems in Fig. 15

applied to NPLTS models. Since a major difference with $\sim_{\text{PTe-tbt}}$ is the use of restricted schedulers, an adaptation of the testing equivalence of [16] to a common model should lead to an equivalence that is coarser than $\sim_{\text{PTe-tbt}}$.

It can however be shown that the two equivalences are different if attention is restricted to a common submodel that does not permit internal nondeterminism. Indeed, absence of internal nondeterminism makes label massaging unnecessary, and we have that reactive probabilistic processes constitute the largest submodel common to the model of [16] and NPLTS. Consider the two reactive probabilistic processes depicted as NPLTS models in Fig. 15, and suppose that what is called synchronization nondeterminism in [16] is handled without using τ inside the labels of the transitions of the interaction systems. The two processes are discriminated by $\sim_{\text{PTe-tbt}}$ because, if we consider the test in the same figure and the maximal resolutions shown in Fig. 16 of the interaction systems, the success probability p_1 of trace $a\,b$ in the second maximal resolution of (s_2, o) is not matched by the success probability 1 of the only maximal resolution of (s_1, o) having a maximal computation labeled with $a\,b$. In contrast, the testing equivalence of [16] cannot distinguish the two processes. Whenever they remain in the interaction system with an arbitrary test, the two identical choices between b and c in the second process must be resolved in the same way by any restricted scheduler that can only yield consistent resolutions. For instance, the only maximal resolutions of (s_2, o) that are consistent among the four shown in Fig. 16 are the first one (choice of b) and the fourth one (choice of c), and their respective success probabilities 1 and 0 are precisely matched by those of the only two maximal resolutions of (s_1, o).

7. Conclusion

In this paper, we have proposed two variants of trace and testing equivalences, respectively denoted by $\sim_{\text{PTr}}$ and $\sim_{\text{PTe-tbt}}$, for the general class of nondeterministic and probabilistic processes, which enjoy desirable properties like:

(1) being preserved by parallel composition,
(2) being fully conservative extensions of the corresponding equivalences studied for nondeterministic processes and for probabilistic processes, and
(3) guaranteeing that trace equivalence is coarser than testing equivalence.

For both equivalences, we have assumed history-independent centralized schedulers. In particular, we have considered the impact of employing deterministic schedulers or randomized schedulers to resolve nondeterminism. We have denoted by $\sim_{\text{PTr}}^{\text{ct}}$ and $\sim_{\text{PTe-tbt}}^{\text{ct}}$ the equivalence variants based on randomized schedulers.

The most studied trace and testing equivalences known in the literature of nondeterministic and probabilistic processes, namely the probabilistic trace-distribution equivalence

$\sim^{\mathrm{ct}}_{\mathrm{PTr,dis}}$ investigated in [32, 7, 27, 28, 6] and the probabilistic testing equivalence $\sim_{\text{PTe-}\sqcup\sqcap}$ investigated in [39, 21, 33, 12], do not fulfill all of these properties. In particular, $\sim^{\mathrm{ct}}_{\mathrm{PTr,dis}}$ is not a congruence with respect to parallel composition and $\sim_{\text{PTe-}\sqcup\sqcap}$ is not a fully conservative extension of the testing equivalences defined in [11] for fully nondeterministic processes, in [9] for generative probabilistic processes, and in [25] for reactive probabilistic processes. Moreover, while the discriminating power of $\sim_{\text{PTe-}\sqcup\sqcap}$ is independent from the use of deterministic of randomized schedulers, the inclusion of this testing equivalence in the trace-distribution equivalence heavily depends on the use of randomized schedulers when defining the trace semantics. Specifically, we have that $\sim_{\text{PTe-}\sqcup\sqcap}$ is contained in $\sim^{\mathrm{ct}}_{\mathrm{PTr,dis}}$ but not in $\sim_{\mathrm{PTr,dis}}$, being the former based on randomized schedulers and the latter on deterministic schedulers.

The main idea behind the new trace equivalence $\sim_{\mathrm{PTr}}$ that we have proposed is that of comparing the execution probabilities of single traces rather than entire trace distributions, so as to avoid debatable distinctions such as the one made by $\sim_{\mathrm{PTr,dis}}$ in Fig. 4. This requires a shift from considering fully matching resolutions to considering partially matching resolutions, which opens the way to compositionality under centralized schedulers.

The main ideas behind the new testing equivalence $\sim_{\text{PTe-tbt}}$ are: (i) matching all resolutions on the basis of their success probabilities, rather than taking into account only maximal and minimal success probabilities, and (ii) considering success probabilities in a trace-by-trace fashion, rather than cumulatively on entire resolutions. It is the trace-by-trace approach that annihilates the impact of the copying capability introduced by observers not of the same nature as the processes under test, and thus permits defining an equivalence that is fully conservative with respect to classical testing equivalences. Remarkably, we have seen in Thm. 5.4 that our new approach, when restricted to fully nondeterministic processes, generative probabilistic processes, and reactive probabilistic processes, yields the same testing equivalences longly studied in the literature.

In order to get to the trace-by-trace approach, it has been important to pass through an additional testing semantics, $\sim_{\text{PTe-}\forall\exists}$, which is not fully backward compatible with testing semantics for restricted classes of processes but, unlike $\sim_{\text{PTe-}\sqcup\sqcap}$, it implies trace semantics. This testing semantics does act as a trait d'union between the testing semantics focussing only on extremal success probabilities – because $\sim^{\mathrm{ct}}_{\text{PTe-}\forall\exists}$ coincides with $\sim_{\text{PTe-}\sqcup\sqcap}$ – and our new fully backward compatible testing semantics comparing success probabilities trace-by-trace – because $\sim_{\text{PTe-}\forall\exists}$ coincides with $\sim_{\text{PTe-tbt,dis}}$.

Another interesting result about testing semantics is that using randomized schedulers to resolve nondeterminism annihilates the difference between many equivalences. Indeed, we have that $\sim^{\mathrm{ct}}_{\text{PTe-tbt,dis}}$ coincides with $\sim^{\mathrm{ct}}_{\text{PTe-}\forall\exists}$ and with $\sim^{\mathrm{ct}}_{\text{PTe-}\sqcup\sqcap}$, which in turn coincides with $\sim_{\text{PTe-}\sqcup\sqcap}$, its variant based on deterministic schedulers. Thus, $\sim^{\mathrm{ct}}_{\text{PTe-tbt,dis}}$ constitutes an alternative characterization of $\sim_{\text{PTe-}\sqcup\sqcap}$, a fact that reconciles the testing equivalence deeply investigated in the literature with the three approaches recently explored in [4] to the definition of behavioral relations for NPLTS models.

We would like to mention that $\sim_{\mathrm{PTr}}$ and $\sim_{\text{PTe-tbt}}$ did pop up when working in the framework of ULTraS [3]. This is a parametric model encompassing many others such as labeled transition systems, discrete-/continuous-time Markov chains, and discrete-/continuous-time Markov decision processes without/with internal nondeterminism. On this unifying model, we have defined trace, testing, and bisimulation equivalences in an abstract way and shown that they induce new equivalences (like $\sim_{\mathrm{PTr}}$ and $\sim_{\text{PTe-tbt}}$) different from those known in the literature (like $\sim^{\mathrm{ct}}_{\mathrm{PTr,dis}}$ and $\sim_{\text{PTe-}\sqcup\sqcap}$) when instantiating the model to the NPLTS case.

In this paper, we have also studied the relationships between our new testing semantics and previously defined failure semantics for nondeterministic and probabilistic processes. While in the fully nondeterministic case the two semantics coincide [10], we have shown that $\sim_{\text{PTe-tbt,dis}}$ is strictly finer than $\sim_{\text{PF,dis}}$, while $\sim_{\text{PTe-tbt}}$ is strictly coarser than $\sim_{\text{PF}}$. We conjecture that the former two equivalences and the latter two equivalences respectively coincide if, in the trace-by-trace approach, we compare not only trace-based probabilities of reaching success, but also failure probabilities, i.e., the probabilities of performing maximal computations compatible with a certain trace that do not reach success.

As future work, we plan to study equational and logical characterizations of the new trace and testing equivalences that we have introduced in this paper.

Acknowledgement

We would like to thank the anonymous referees for their stimulating comments and Marco Tinacci for his useful suggestions on the comparison with [16]. This work has been partially supported by the FP7-IST-FET Project ASCENS, grant no. 257414, by the EU Project QUANTICOL, grant no. 600708, and by the MIUR project CINA.

References

[1] S. Abramsky. Observational equivalence as a testing equivalence. *Theoretical Computer Science*, 53:225–241, 1987.

[2] M. Bernardo, R. De Nicola, and M. Loreti. Revisiting trace and testing equivalences for nondeterministic and probabilistic processes. In *Proc. of the 15th Int. Conf. on Foundations of Software Science and Computation Structures (FOSSACS 2012)*, volume 7213 of *LNCS*, pages 195–209. Springer, 2012.

[3] M. Bernardo, R. De Nicola, and M. Loreti. A uniform framework for modeling nondeterministic, probabilistic, stochastic, or mixed processes and their behavioral equivalences. *Information and Computation*, 225:29–82, 2013.

[4] M. Bernardo, R. De Nicola, and M. Loreti. Relating strong behavioral equivalences for processes with nondeterminism and probabilities. *Theoretical Computer Science*, 2014. To appear.

[5] S.D. Brookes, C.A.R. Hoare, and A.W. Roscoe. A theory of communicating sequential processes. *Journal of the ACM*, 31:560–599, 1984.

[6] L. Cheung, N. Lynch, R. Segala, and F. Vaandrager. Switched PIOA: Parallel composition via distributed scheduling. *Theoretical Computer Science*, 365:83–108, 2006.

[7] L. Cheung, M. Stoelinga, and F. Vaandrager. A testing scenario for probabilistic processes. *Journal of the ACM*, 54(6), 2007.

[8] I. Christoff. Testing equivalences and fully abstract models for probabilistic processes. In *Proc. of the 1st Int. Conf. on Concurrency Theory (CONCUR 1990)*, volume 458 of *LNCS*, pages 126–140. Springer, 1990.

[9] R. Cleaveland, Z. Dayar, S.A. Smolka, and S. Yuen. Testing preorders for probabilistic processes. *Information and Computation*, 154:93–148, 1999.

[10] R. De Nicola. Extensional equivalences for transition systems. *Acta Informatica*, 24:211–237, 1987.

[11] R. De Nicola and M. Hennessy. Testing equivalences for processes. *Theoretical Computer Science*, 34:83–133, 1984.

[12] Y. Deng, R.J. van Glabbeek, M. Hennessy, and C. Morgan. Characterising testing preorders for finite probabilistic processes. *Logical Methods in Computer Science*, 4(4:4):1–33, 2008.

[13] Y. Deng, R.J. van Glabbeek, M. Hennessy, C. Morgan, and C. Zhang. Remarks on testing probabilistic processes. In *Computation, Meaning, and Logic: Articles Dedicated to Gordon Plotkin*, volume 172 of *ENTCS*, pages 359–397. Elsevier, 2007.

[14] Y. Deng, R.J. van Glabbeek, C. Morgan, and C. Zhang. Scalar outcomes suffice for finitary probabilistic testing. In *Proc. of the 16th European Symp. on Programming (ESOP 2007)*, volume 4421 of *LNCS*, pages 363–378. Springer, 2007.

[15] C. Derman. *Finite State Markovian Decision Processes*. Academic Press, 1970.
[16] S. Georgievska and S. Andova. Probabilistic may/must testing: Retaining probabilities by restricted schedulers. *Formal Aspects of Computing*, 24:727–748, 2012.
[17] A. Giacalone, C.-C. Jou, and S.A. Smolka. Algebraic reasoning for probabilistic concurrent systems. In *Proc. of the 1st IFIP Working Conf. on Programming Concepts and Methods (PROCOMET 1990)*, pages 443–458. North-Holland, 1990.
[18] H. Hansson and B. Jonsson. A calculus for communicating systems with time and probabilities. In *Proc. of the 11th IEEE Real-Time Systems Symp. (RTSS 1990)*, pages 278–287. IEEE-CS Press, 1990.
[19] H. Jifeng, K. Seidel, and A. McIver. Probabilistic models for the guarded command language. *Science of Computer Programming*, 28:171–192, 1997.
[20] B. Jonsson, C. Ho-Stuart, and W. Yi. Testing and refinement for nondeterministic and probabilistic processes. In *Proc. of the 3rd Int. Symp. on Formal Techniques in Real Time and Fault Tolerant Systems (FTRTFT 1994)*, volume 863 of *LNCS*, pages 418–430. Springer, 1994.
[21] B. Jonsson and W. Yi. Compositional testing preorders for probabilistic processes. In *Proc. of the 10th IEEE Symp. on Logic in Computer Science (LICS 1995)*, pages 431–441. IEEE-CS Press, 1995.
[22] B. Jonsson and W. Yi. Testing preorders for probabilistic processes can be characterized by simulations. *Theoretical Computer Science*, 282:33–51, 2002.
[23] C.-C. Jou and S.A. Smolka. Equivalences, congruences, and complete axiomatizations for probabilistic processes. In *Proc. of the 1st Int. Conf. on Concurrency Theory (CONCUR 1990)*, volume 458 of *LNCS*, pages 367–383. Springer, 1990.
[24] R.M. Keller. Formal verification of parallel programs. *Communications of the ACM*, 19:371–384, 1976.
[25] M. Kwiatkowska and G. Norman. A testing equivalence for reactive probabilistic processes. In *Proc. of the 5th Int. Workshop on Expressiveness in Concurrency (EXPRESS 1998)*, volume 16(2) of *ENTCS*, pages 114–132. Elsevier, 1998.
[26] K.G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94:1–28, 1991.
[27] N.A. Lynch, R. Segala, and F. Vaandrager. Compositionality for probabilistic automata. In *Proc. of the 14th Int. Conf. on Concurrency Theory (CONCUR 2003)*, volume 2761 of *LNCS*, pages 208–221. Springer, 2003.
[28] A. Parma and R. Segala. Axiomatization of trace semantics for stochastic nondeterministic processes. In *Proc. of the 1st Int. Conf. on the Quantitative Evaluation of Systems (QEST 2004)*, pages 294–303. IEEE-CS Press, 2004.
[29] A. Philippou, I. Lee, and O. Sokolsky. Weak bisimulation for probabilistic systems. In *Proc. of the 11th Int. Conf. on Concurrency Theory (CONCUR 2000)*, volume 1877 of *LNCS*, pages 334–349. Springer, 2000.
[30] M.O. Rabin. Probabilistic automata. *Information and Control*, 6:230–245, 1963.
[31] R. Segala. *Modeling and Verification of Randomized Distributed Real-Time Systems*. PhD Thesis, 1995.
[32] R. Segala. A compositional trace-based semantics for probabilistic automata. In *Proc. of the 6th Int. Conf. on Concurrency Theory (CONCUR 1995)*, volume 962 of *LNCS*, pages 234–248. Springer, 1995.
[33] R. Segala. Testing probabilistic automata. In *Proc. of the 7th Int. Conf. on Concurrency Theory (CONCUR 1996)*, volume 1119 of *LNCS*, pages 299–314. Springer, 1996.
[34] R. Segala and N.A. Lynch. Probabilistic simulations for probabilistic processes. In *Proc. of the 5th Int. Conf. on Concurrency Theory (CONCUR 1994)*, volume 836 of *LNCS*, pages 481–496. Springer, 1994.
[35] K. Seidel. Probabilistic communicating processes. *Theoretical Computer Science*, 152:219–249, 1995.
[36] A. Sokolova and E.P. de Vink. Probabilistic automata: System types, parallel composition and comparison. In *Validation of Stochastic Systems*, volume 2925 of *LNCS*, pages 1–43. Springer, 2004.
[37] R.J. van Glabbeek, S.A. Smolka, and B. Steffen. Reactive, generative and stratified models of probabilistic processes. *Information and Computation*, 121:59–80, 1995.
[38] M.Y. Vardi. Automatic verification of probabilistic concurrent finite-state programs. In *Proc. of the 26th IEEE Symp. on Foundations of Computer Science (FOCS 1985)*, pages 327–338. IEEE-CS Press, 1985.
[39] W. Yi and K.G. Larsen. Testing probabilistic and nondeterministic processes. In *Proc. of the 12th Int. Symp. on Protocol Specification, Testing and Verification (PSTV 1992)*, pages 47–61. North-Holland, 1992.